UNIVERSIDADE DE SÃO PAULO
Instituto de Ciências Matemáticas e de Computação

Beyond the comfort zone: digital nudges for off-profile recommendations

Gabrielle Aparecida Pires Alves

Dissertação de Mestrado do Programa de Pós-Graduação em Ciências de Computação e Matemática Computacional (PPG-CCMC)

ICMC USP
SÃO CARLOS

SERVIÇO DE PÓS-GRADUAÇÃO DO ICMC-USP

Data de Depósito:

Assinatura: ______________________

Gabrielle Aparecida Pires Alves

Beyond the comfort zone: digital nudges for off-profile recommendations

Master dissertation submitted to the Instituto de Ciências Matemáticas e de Computação – ICMC-USP, in partial fulfillment of the requirements for the degree of the Master Program in Computer Science and Computational Mathematics. *EXAMINATION BOARD PRESENTATION COPY*

Concentration Area: Computer Science and Computational Mathematics

Advisor: Prof. Dr. Marcelo Garcia Manzato

USP – São Carlos
April 2023

Ficha catalográfica elaborada pela Biblioteca Prof. Achille Bassi
e Seção Técnica de Informática, ICMC/USP,
com os dados inseridos pelo(a) autor(a)

P113b

Pires Alves, Gabrielle Aparecida
Beyond the comfort zone: digital nudges for off-profile recommendations / Gabrielle Aparecida Pires Alves; orientador Marcelo Garcia Manzato. -- São Carlos, 2023.
89 p.

Dissertação (Mestrado - Programa de Pós-Graduação em Ciências de Computação e Matemática Computacional) -- Instituto de Ciências Matemáticas e de Computação, Universidade de São Paulo, 2023.

1. discovery. 2. nudging. 3. recommender system. 4. diversity. 5. bias. I. Garcia Manzato, Marcelo, orient. II. Título.

Bibliotecários responsáveis pela estrutura de catalogação da publicação de acordo com a AACR2:
Gláucia Maria Saia Cristianini - CRB - 8/4938
Juliana de Souza Moraes - CRB - 8/6176

Gabrielle Aparecida Pires Alves

Além da zona de conforto: nudges digitais para recomendações fora de perfil

Dissertação apresentada ao Instituto de Ciências Matemáticas e de Computação – ICMC-USP, como parte dos requisitos para obtenção do título de Mestre em Ciências – Ciências de Computação e Matemática Computacional. *EXEMPLAR DE DEFESA*

Área de Concentração: Ciências de Computação e Matemática Computacional

Orientador: Prof. Dr. Marcelo Garcia Manzato

USP – São Carlos
Abril de 2023

Este trabalho é dedicado às crianças adultas que,
quando pequenas, sonharam em ajudar a melhorar o mundo.

ACKNOWLEDGEMENTS

Quero expressar minha gratidão à minha família e amigos, que estiveram comigo antes e durante deste mestrado, me ensinando e apoiando em todos os momentos.

Gostaria de agradecer ao Professor Manzato por ter apoiado todas as minhas loucuras e projetos, e ao Professor Jannach por sua mentoria e paciência. Ambos me ensinaram muito sobre como bons orientadores são moldados.

Também gostaria de agradecer à Universidade de São Paulo e ao Instituto de Ciências Matemáticas e de Computação, por me inspirarem a me apaixonar pela pesquisa.

À Coordenação de Aperfeiçoamento de Pessoal de Nível Superior (CAPES) - Código de Financiamento 001 - pelo apoio financeiro que tornou possível a minha jornada acadêmica.

"As invenções são, sobretudo,
o resultado de um trabalho de teimoso."
(Santos Dumont)

RESUMO

ALVES, G. **Além da zona de conforto: nudges digitais para recomendações fora de perfil** . 2023. 89 p. Dissertação (Mestrado em Ciências – Ciências de Computação e Matemática Computacional) – Instituto de Ciências Matemáticas e de Computação, Universidade de São Paulo, São Carlos – SP, 2023.

Na sociedade em constante evolução de hoje, os sistemas de recomendação surgiram como uma ferramenta crucial na era digital. Em muitas áreas de aplicação desses sistemas, como em sites de streaming de mídia, o objetivo principal dos provedores de serviço de recomendação é aumentar o engajamento dos usuários ajudando-os a descobrir novos tipos de conteúdo que gostem. Algoritmos tradicionais de filtragem colaborativa geralmente levam a algum nível de descoberta. No entanto, em alguns casos, é útil promover conteúdo fora do perfil para que os usuários possam explorar as possibilidades. Mas, ao mostrar itens fora do perfil para os usuários junto com itens familiares, é possível que os novos itens passem despercebidos. Nesta pesquisa, analisamos como a técnica de "digital nudging", ou seja, o estímulo sutil das escolhas do usuário em uma direção específica, pode ajudar a aumentar a atenção e o interesse dos usuários em conteúdo fora do seu perfil. Fizemos um estudo com usuários (N=1.064) em um aplicativo real de recomendação de livros. Descobrimos que, ao direcionar os usuários para livros recomendados de gêneros não preferidos, houve um aumento significativo de livros fora do perfil adicionados às listas de leitura, confirmando a eficácia do digital nudging. No entanto, também percebemos que a percepção subjetiva da relevância das recomendações pode diminuir quando usamos nudging, resultando em menor satisfação e menor intenção de reutilizar o sistema. Portanto, o digital nudging em recomendações é eficaz a curto prazo, mas deve ser usado com cautela, mantendo-se atento às percepções gerais de qualidade dos usuários e aos potenciais efeitos prejudiciais a longo prazo.

Palavras-chave: descoberta, nudging, sistemas de recomendação, diversidade, vies.

ABSTRACT

ALVES, G. **Beyond the comfort zone: digital nudges for off-profile recommendations**. 2023. 89 p. Dissertação (Mestrado em Ciências – Ciências de Computação e Matemática Computacional) – Instituto de Ciências Matemáticas e de Computação, Universidade de São Paulo, São Carlos – SP, 2023.

In today's constantly evolving society, recommendation systems have emerged as a crucial tool in the digital age. In many application areas of recommendation systems, such as media streaming sites, the main objective of service providers is to increase user engagement by helping them discover new types of content they like. Standard collaborative filtering algorithms typically lead to some level of discovery. However, in some cases, it is useful to actively promote out-of-profile items so that users can explore new possibilities. But when showing out-of-profile items to users along with more familiar ones, new items may go unnoticed. In this research, we examine how the technique of "digital nudging", or subtly directing user choices in a specific direction, can help increase user attention and interest in out-of-profile content. We conducted a user study (N=1,064) in a real book recommendation application. We found that by directing users to recommended books of non-preferred genres, there was a significant increase in out-of-profile books added to reading lists, confirming the effectiveness of digital nudging. However, we also found that the subjective perception of recommendation relevance may decrease when using nudging, resulting in lower satisfaction and lower intention to reuse the system. Therefore, digital nudging in recommendations is effective in the short term, but should be used with caution, keeping an eye on users' overall perceptions of quality and potential long-term harmful effects.

Keywords: discovery, nudging, recommender system, diversity, bias.

LIST OF FIGURES

LIST OF TABLES

LIST OF ABBREVIATIONS AND ACRONYMS

CFI	comparative fit index
ILS	intra-list similarity
MVVM	Model-View-View Model
ResQue	Recommender systems' Quality of user experience
RMSEA	Root Mean Square Error of Approximation
RS	recommender system

CONTENTS

CHAPTER

1

INTRODUCTION

"It's only after you've stepped outside your comfort zone that you begin to change, grow, and transform."

Roy T. Bennett

In recent years, the amount of information created by society has grown to an unprecedented scale, and our brains are not adequately equipped to cope with it. This has led to a condition known as information overload (AKIN, 1997; JACOBY; SPELLER; BERNING, 1974), which occurs when the amount of information provided surpasses an individual's ability to process it. As a consequence, this can result in stress, confusion, and anxiety, ultimately impairing one's ability to prioritize and make decisions (MALHOTRA, 1982). Interestingly, an individual's performance and the quality of their decisions are positively correlated with the amount of information they receive up to a certain point. However, information overload sets in once that threshold is surpassed, and their performance rapidly declines (EPPLER; MENGIS, 2004).

Any situation that requires decision-making, such as browsing the internet, selecting a movie to watch, or buying a new book, can result in information overload. To combat this challenge, strategies have been developed to manage information and its various characteristics, including quantity, frequency, intensity, and quality. One such strategy involves using recommender system (RS), which can assist individuals in making decisions. These systems automate some of the decision-making processes and aim to offer affordable, customized, and high-quality recommendations (JANNACH *et al.*, 2010).

As the world becomes increasingly digital and interactive, recommender systems are a powerful tool to help us navigate through the vast amount of information available and find items that we truly need (RESNICK; VARIAN, 1997; GE *et al.*, 2012). These systems personalize recommendations based on users' interests, preferences, and behavior, providing an enhanced user experience (CHEN *et al.*, 2020). To enable effective personalization in recommender systems, the use of a user model or profile is crucial. Such models can be built using a variety of techniques, including collaborative, content-based, knowledge-based, and hybrid approaches.

Collaborative systems rely on users' past behaviors to predict their future interests. Content-based recommendation systems use item descriptions and assign importance to their characteristics to make recommendations. Knowledge-based systems utilize additional, often manually provided, information about the user and the available items to make recommendations. Finally, hybrid approaches combine different techniques to generate better and more precise recommendations, taking into account the advantages and disadvantages of each technique (JANNACH *et al.*, 2010). Aside from that, a number of novel recommendation approaches have also been proposed, including social network-based recommender systems, fuzzy recommender systems, context-aware recommender systems, and group recommender systems (LU *et al.*, 2015).

The development of these recommender systems' techniques has led to their widespread adoption in real-world applications. However, while recommendation systems have many benefits in the digital world, it is important to recognize that personalizing our experience based on our taste and personality traits can have potentially negative consequences. Jannach *et al.* (2010) note that we must consider the broader impact of recommender systems, which cannot be reduced to basic decision theory principles. Therefore, it is crucial to explore different approaches and understand their implications on both individuals and society at large, given that they influence user preferences and steer decision-making, at both individual and collective levels (MILANO; TADDEO; FLORIDI, 2020).

Recommender systems have become an integral part of our digital lives, shaping our experiences and interactions. However, their commercial-driven design and deployment have raised ethical concerns regarding privacy, autonomy, and the well-being of individuals (MILANO; TADDEO; FLORIDI, 2020). One of the primary ethical challenges caused by recommender systems is the creation of filter bubbles that limit users' exposure to diverse perspectives and content. This narrow range of options can lead to unfair outcomes and harm the autonomy and personal identity of the users (CHEN *et al.*, 2020).

Filter bubbles, or the limited exposure to diverse views and content (BAEZA-YATES, 2018), occur when a RS only suggests items that closely relate to topics of interest in the user's profile, resulting in a narrow range of options and preventing users from discovering new and high-quality items. This leads to a situation where individuals are insulated from diverse perspectives, leading to the creation of an unfair and unbalanced environment. In response, the RS community has shifted its focus beyond accuracy, to address the limitations and biases inherent in the current approach of collaborative filtering.

To mitigate the impact of filter bubbles and biases in RSs, it is essential to adopt a more comprehensive and balanced approach, that not only considers the accuracy of recommendations but also the diversity and fairness of the results (LI *et al.*, 2017; SCHELENZ, 2021). This shift in focus is critical to ensure that RSs serve as a tool for enriching individuals' lives and promoting a fairer and more diverse society.

This research project aims to address the explore questions of discovery and exploration in recommender systems. While they are essential in practice, there is comparably little research in the academic literature that explicitly aims at understanding how discovery support affects the user experience and quality perception of users. By examining various aspects of the user experience, we aim to engage users, helping them break out of their comfort zone and exposing them to a broader spectrum of information and perspectives, while preserving the quality of the system.

1.1 Motivation and Research Questions

Nowadays, recommender systems are widely used in online services such as Amazon, Netflix, Spotify, and YouTube to provide personalized recommendations to users. While these systems aim to match users with items they are likely to enjoy based on their past preferences, they also strive to suggest items that are outside their historical tastes to support user discovery. However, there is a lack of research in the academic literature that specifically examines how discovery support affects the user experience and perception of quality. Some studies have investigated discovery in the music domain, such as (CELMA, 2020) and (LUDEWIG *et al.*, 2021), but these are limited in scope. Most studies have been based on data analysis and offline experimentation, with a focus on the role of diversity and novelty in recommender systems. However, it is unclear how well metrics such as intra-list similarity (ILS) reflect users' diversity perceptions.

A significant question that arises is whether recommending "off-profile" items, i.e., items outside a user's past preferences that aim to promote discovery, would affect users' perception of quality and behavior. For instance, Ekstrand et al. (EKSTRAND *et al.*, 2014) reported that suggesting surprising or unexpected items may lead to negative impacts on user satisfaction with recommendations. However, other studies such as Ge et al. (GE; GEDIKLI; JANNACH, 2011; GE *et al.*, 2012) indicate that adding unsuitable but diverse items, such as recommending comedies to users who prefer action movies, may not be noticed by users, depending on the item's position in the recommended list.

This study aims to investigate the impact of subtly guiding users towards off-profile items that they may not have noticed. In order to achieve this, we implemented a technique known as nudge. Nudging is a concept introduced by Thaler and Sunstein (2009), which involves gently guiding people towards a desired behavior without limiting their options. In the digital world, nudging can also be employed to guide users towards certain options or behaviors, often referred to as digital nudging. In this study, digital nudges will be used to encourage users to consider items outside of their typical preferences.

In light of this context, the following research questions are proposed to structure the motivating statement of the problem for this research:

- How does digitally nudging users to consider off-profile items impact their choice behavior?
- How does digitally nudging users to consider off-profile items impact their quality perception of recommended items?

The goal of this research is to investigate how *digitally nudging* (CARABAN *et al.*, 2019; SCHNEIDER; WEINMANN; BROCKE, 2018) towards off-profile items affects their behavior and perception of recommended items. The study will specifically focus on two aspects: (i) the actual choices made by users, and (ii) their perception of the quality of recommendations.

1.2 Contributions

This study focuses on investigating the questions mentioned above in the context of book recommendations. We conducted an online study with 1,064 users of a social book recommendation site. Participants were presented with recommendation lists that included both books matching their preferred genres and off-profile recommendations from other genres. The treatment group received different types of nudges on the off-profile items to expand their range of interests and discover books in new genres. It is worth noting that unlike previous research such as (BERGER; MÜLLER; NÜSKE, 2020; JESSE; JANNACH; GULA, 2021), our goal is not to determine the most effective type of nudge in a particular domain. Rather, we aim to analyze the combined effects of the applied nudges on user behavior.

According to our study, nudging was highly effective in encouraging participants to add off-profile items to their reading lists. However, some of the nudged participants perceived the recommendations as less relevant to their interests and were less satisfied with their reading lists, even though the options were the same as those in the control group. This finding suggests that there are interesting psychological phenomena that arise when off-profile items receive more attention through nudging.

In conclusion, we found that nudging was effective in this domain, at least in the short-term. However, we also found that digital nudges must be designed with caution to avoid potential negative long-term effects, such as a lower intention to rely on the recommendations in the future. Therefore, it is important to carefully consider the design and implementation of digital nudges to ensure that they achieve the desired outcomes without unintended consequences.

1.3 Master's Disseration Outline

The master's dissertation is organized as follows. In Chapter 2, we provide background information on recommender systems, biases, diversity, and nudges. In Chapter 3, we present related works. Chapter 4 describes the user study, and Chapter 5 details the analysis and interpretation of the results. Finally, we discuss research limitations and potential future works.

CHAPTER

2

BACKGROUND

To create a system that serves the varying needs of users, we must examine recommender systems, biases, and nudge theory. This chapter explores relevant literature, highlighting advancements in these areas. By utilizing these concepts, we can design a system that not only recommends content, but also encourages unbiased choices that align with individual preferences.

2.1 Recommender System

The internet's vast expansion has led to an abundance of data, making it challenging for users to find the information they seek (JANNACH *et al.*, 2010). To address this, Recommendation Systems (RSs) emerged in the 1990s, aiming to provide personalized recommendations that match user preferences.

At the most basic level, RSs offer ranked lists of items that are most relevant to the user's profile (RICCI; ROKACH; SHAPIRA, 2011). By utilizing various techniques, such as explicit and implicit approaches, these systems aim to filter and recommend products or services tailored to the user's needs and preferences, improving their overall experience and saving them time and effort.

The Netflix Prize in 2006, which aimed to enhance the accuracy of RSs, led to the development of new methodologies and improved existing algorithms, further advancing the research in this field. In the following sections, we will discuss some of the techniques used in modern recommendation systems.

2.1.1 Classical Approaches

Recommendation systems (RSs) are built on three fundamental components: items, users, and their interactions. Each of these components plays a crucial role in the functioning of an RS (RICCI; ROKACH; SHAPIRA, 2011).

Item. Items refer to the objects that are recommended and can be associated with a cognitive or monetary cost. They can be classified as positive (useful) or negative (not relevant) based on their relevance to the user.

User. Users have distinct characteristics, and RSs utilize various user information to personalize recommendations and enhance human-computer interaction. The information can be organized in different ways, and the selection of what information to use depends on the recommendation technique.

Interactions. Interactions refer to the logs that are generated during the interaction between the system and the user. These logs contain information such as the item chosen by the user and the context of the interaction, and they are critical to understanding the user's preferences and improving the recommendation process.

Designing more effective RSs that offer personalized and relevant recommendations to users can be achieved by comprehending the roles and characteristics of the three key components of RSs. These components are present in all RSs and different approaches can be employed to achieve personalized recommendations. Based on a taxonomy (BURKE; ROBIN, 2007), the various recommendation approaches can be categorized into six distinct classes:

Collaborative filtering. Collaborative filtering is the most popular and widespread approach in recommendation systems. It operates under the assumption that users who have shared interests in the past will also have similar preferences in the future, making it possible to predict what the user might like or be interested in based on their behavior. Collaborative filtering algorithms can be classified into two categories: Collaborative Filtering Based on Items, which seeks to find a similarity metric between items, but it can present problems related to sparsity due to items that have not been evaluated by users, for example; and Collaborative Filtering Based on Users, which uses filtering based on a neighborhood to identify other users with similar preferences in the past and create a prediction over items not yet viewed by the user based on a set of users with similar preferences.

Content-based. In contrast to collaborative recommendation, which does not require item details, content-based recommendation considers the item's characteristics. To generate recommendations, metadata and explicit information about the items are assimilated with the users' preferences. One of the most popular methods used in RSs for this approach is K-Nearest Neighbor (KNN) algorithms. When using content-based recommendations, past items liked by the user and metadata such as genres are used to predict the next recommendation.

Demographic. In the demographic approach, recommendations are tailored to the user's profile, such as their language or country. This type of personalization is commonly used on websites that serve users from different regions or with different language preferences.

Knowledge-based. Knowledge-based systems recommend items based on domain knowledge, which assesses how well the item meets users' needs and preferences, and how useful the item is for the user. To implement this approach, the system first measures the user's needs, which is usually represented as a problem description. Then, it matches the problem description with the recommendations, which are solutions to the problem. The system measures how closely the user's needs match the recommendations, and uses this information to generate personalized recommendations.

Community-based. This approach has gained relevance and popularity with the growth of social networks. This technique uses information about the users and their social connections, or users in the same community, to create personalized recommendations. By leveraging the social connections of users, this approach aims to identify items that are likely to be of interest to a user based on the interests and preferences of their friends or peers. This approach has the potential to provide more accurate and relevant recommendations, as users tend to share similar interests and preferences with their social connections.

Hybrid recommender systems. To make the most effective recommendations, different approaches can be used to exploit information in various ways. Each approach has its positive and negative points. The collaborative approach requires a large amount of data, while the content-based approach focuses on the content itself, even if it's limited, to make a recommendation. However, the hybrid recommendation approach combines features of different approaches to create more accurate and personalized predictions.

Recommendation systems seek to provide data-driven, personalized solutions by filtering information. However, it is important to acknowledge that these systems can create biases. In the next subsection, we will explore how biases can arise in recommendation systems.

2.1.2 Biases and Filter Bubbles

Recommendation Systems have had a significant impact on various applications; however, they face issues related to bias that can challenge their effectiveness (ELSWEILER; TRATTNER; HARVEY, 2017). The fundamental concept behind RS is to personalize the user's experience by filtering data and developing solutions that fit their profile. However, this approach can be problematic as it exposes users to a limited number of items and restricts interactions, which may result in the failure to offer more diverse options that could also match the user's preferences. This could be classified as unfairness (CHEN *et al.*, 2020).

There are different types of biases, such as:

Selection Bias is a type of bias that occurs when users have the freedom to select which items to rate, leading them to rate only the items they like or only particularly good or bad items. This can result in a skewed dataset that fails to accurately represent the user's preferences and can impact the recommendations provided by the RS.

Conformity Bias is another type of bias that can distort users' feedback and judgment. This bias happens when users tend to behave similarly to others in their group, even if their true preferences are different. This behavior can cause users to rate items that they may not have liked or rate items higher or lower than they would have on their own. This can result in less diverse recommendations and a limited range of options for users.

Exposure Bias is another type of bias that can occur in recommendation systems. This happens when users are only exposed to a portion of the feedback about an item, which can cause the overall feedback about the item, including both positive and negative aspects, to go unnoticed. As a result, the recommendations can be skewed towards only popular or positively-rated items, and users may miss out on items that may be a better fit for them.

Position Bias is a widely recognized type of bias in recommendation systems. This bias occurs when users interact more frequently with items that are placed higher up on the list, regardless of whether those items are actually relevant to their preferences or not. For example, a user may select an item from the top of the list because they assume it is the most popular or highly recommended, even if it does not align with their personal preferences.

Inductive Bias is a type of bias in machine learning models that refers to the tendency of the model to make inferences and generalize beyond the training data it was provided. In other words, the model may learn patterns and assumptions from the training data that do not hold true for the entire population. This can lead to inaccuracies in predictions when applied to new and unseen data.

Popularity Bias is a type of bias in recommendation systems where popular items are recommended more frequently, leading to a decreased level of personalization and serendipity, which can ultimately negatively affect the user experience. This bias can exacerbate the popularity of already popular items and overshadow the recommendation of items that are better suited to the user's preferences.

Unfairness recommendations are challenging to define due to subjectivity and varying contexts. Wang *et al.* (2022) note that unfairness can impact both users and items in recommendation scenarios. For instance, less accurate movie and music recommendations are given to female and older users, and diversity and novelty can vary among user groups.

While biased recommendations are acknowledged by (DELDJOO *et al.*, 2022), they do not equate bias with unfairness. Numerous definitions of fairness have been proposed in various disciplines, as it is a societal construct. Fairness is a complex concept with multiple perspectives, and social constructs may vary depending on environmental factors. For this work, we examine unfairness in movie recommendations, where certain items are favored over others due to aspects such as gender, popularity, and budget.

Figure 1 – Relations Between Biases and Unfairness

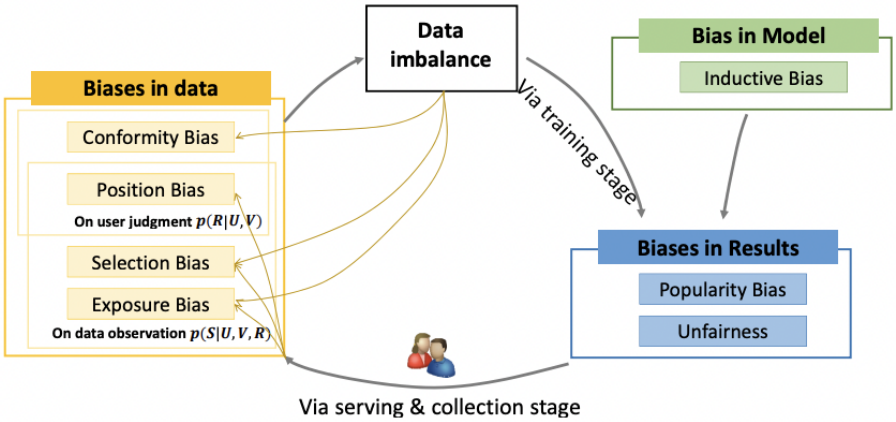

Source: (CHEN *et al.*, 2020)

Exposure to bias and unfairness in recommendation systems is often a result of algorithms that determine which items to show based on factors such as friends, geolocation, and other variables (CHEN *et al.*, 2020). These biases are pervasive in our everyday lives, whether in the digital realm or beyond. Users are influenced by external factors and public opinions, and these same tendencies persist in the online world.

Bias can manifest at various stages of a recommendation system. Figure 1 illustrates how unaddressed biases can become amplified over time, leading to a loop of biases at different layers of the recommendation system. Collectively, these biases contribute to filter bubbles that systematically create unfairness and harm individuals or groups in favor of others (CHEN *et al.*, 2020).

The term "filter bubble" refers to a phenomenon where personalized information can effectively isolate individuals from diverse viewpoints and content, according to Nguyen *et al.* (2014). Filter bubbles can polarize users and create groups that share the same views, resulting in intellectual isolation and distancing from information that disagrees with their views (GELFERT, 2018). The first indications of this condition appeared in 2009 when companies such as Google began to personalize their services to enhance users' experiences. Users' previous interactions with the system, as well as their explicit preferences, affect their experience with the results (PARISER, 2011).

However, as humans, we inherently possess the ability to learn and grow with new information. A recommendation system with a filter bubble environment inhibits creativity and learning, hindering our natural potential (PARISER, 2011). Therefore, it is crucial to acknowledge the negative impact of filter bubbles on users and promote a diverse range of viewpoints and content. This will encourage growth and adaptation to new information, facilitating our inherent nature to learn and evolve.

Understanding how algorithms contribute to the formation of filter bubbles is essential in addressing this phenomenon. While algorithms are not the sole cause, they do play a significant role in creating filter bubbles (PASSE; DRAKE; MAYGER, 2018). Users tend to consume content that aligns with their opinions and avoid content that challenges them. They also establish connections with those who share their beliefs, which promotes confirmation bias, segregation, and polarization, ultimately leading to filter bubbles (VICARIO *et al.*, 2016).

Filter bubbles present a serious obstacle to democracy and society's ability to understand each other. They are invisible, and people often do not realize they are seeing things differently from others (LUNARDI *et al.*, 2020), which reinforces their perspective as the only valid one. Resolving these issues is challenging, but awareness and diversifying content consumption are crucial steps in combating filter bubbles.

Various strategies have been proposed to tackle filter bubbles, including creating diversity-focused algorithms and using interactive visualizations to break them (MUNSON; RESNICK, 2010). To achieve this, we must go beyond prioritizing accuracy and consider additional quality features to improve recommendations and decrease filter bubbles.

In order to combat the negative impact of filter bubbles on society, it is crucial to understand their mechanisms and implement measures that promote diverse content consumption. Despite the goal of recommendation systems being to provide useful items to the user, some researchers mistakenly prioritize high accuracy (KONSTAN; RIEDL, 2011), perpetuating the persistence of filter bubbles in RS. However, to effectively eliminate filter bubbles, it is essential to prioritize user experience beyond accuracy. Lunardi *et al.* (2020) suggest that considering additional quality features beyond accuracy measures can improve recommendations and decrease the impact of filter bubbles.

2.1.3 Beyond Accuracy Measures

To create a recommendation system that truly meets the needs of users, it is essential to take a holistic approach that considers multiple factors. Beyond just diversity, a range of measures related to user experience should be evaluated in order to create a well-rounded approach. Additionally, it is crucial to incorporate user-centered evaluation metrics when designing experiments for recommendation algorithms. These metrics can help ensure that the system is meeting users' needs and expectations. By taking a comprehensive approach that takes all these factors into account, it is possible to develop a recommendation system that is truly effective in enhancing the overall user experience (RICCI *et al.*, 2010). Next, we will explore some potential user-centered evaluation metrics that can help guide the development and refinement of recommendation systems.

Trust plays a crucial role to assess the effectiveness of a recommendation system. Trust can be defined as the user's confidence in the system's recommendations, which is based on their willingness to rely on the recommendations to achieve positive outcomes (CHOPRA; WALLACE, 2003). To evaluate user trust, it is essential to gather online feedback from users regarding the system's recommendations and their reasonability. Additionally, the number of recommendations followed in an online experiment can be used as an indicator of trust in the RS. This implies that a higher level of trust in the system would result in more recommendations being utilized by the user, ultimately contributing to a more successful recommendation system.

Novelty refers to the suggestion of items that the user is not familiar with (KONSTAN *et al.*, 2006). Measuring the novelty of recommended items can be achieved through online experiments by asking participants whether they are already familiar with them. In a diverse environment, it is recommended to focus on novelty among relevant items only. Additionally, it is important to consider the impact of popularity bias on novelty, as highly popular items are generally assumed to be less novel.

Serendipity refers to the level of unexpectedness in the recommended items. While random recommendations can bring some surprise, the true value of serendipity lies in the amount of relevant and surprising information that the system provides to the user.

Satisfaction corresponds to the feeling of happiness that arises when the user's needs or expectations are met (NANOU; LEKAKOS; FOUSKAS, 2010). This concept is often evaluated alongside other factors, such as trust, in RSs. When a recommendation system enhances the user experience through relevant recommendations and a well-designed human-computer interface, it can lead to user satisfaction. The combination of effective recommendations and a usable interface can increase the user's subjective evaluation of the system. In addition, a system that provides clear explanations for its recommendations may also contribute to user satisfaction.

Diversity is often seen as the antithesis of similarity. Measuring diversity usually involves examining the similarity between items, which is frequently based on content. However, it's important to note that diversity may have an impact on other properties, such as accuracy, which could potentially be compromised as a result.

Effectiveness refers to the ability of the system to recommend items that are relevant and useful to the user and assist the user in making appropriate decisions. While accuracy is typically the key factor that determines the effectiveness of the algorithm, it is also important to consider its potential to introduce the user to new and diverse domains.

Persuasiveness is the ability of a system to influence the behavior of a user and their decision-making. This metric measures the system's capacity to prompt the user to accept a system or make a decision. An example of this persuasive phenomenon is when a user is convinced to purchase a product they previously had no interest in.

Efficiency is an important aspect of a system's performance, which enables users to make quick and informed decisions regarding the most suitable item for their needs. Enhancing efficiency involves comparing the various attributes and qualities of different items and identifying their respective strengths and weaknesses.

In order to counteract the biases present in RSs and prevent the growth of filter bubbles, it is crucial to consider all of the previously mentioned properties. Doing so not only improves the user experience but also encourages users to engage with items outside of their typical preferences. To address the problem of biases and unfairness in recommendation systems, one possible solution is to increase the diversity of item suggestions and recommend items with characteristics that differ from those typically preferred by users, while ensuring that the recommended items remain relevant to the user.

Diversity in recommendations offers novelty and serendipity, both of which directly impact the user experience, trust, and satisfaction with the system. These factors can have a positive or negative effect depending on how relevant the recommendations are.

Developing a system that caters to user preferences is critical for those involved in RSs. Analyzing the factors that influence user choices can prove instrumental in identifying areas for improvement in future projects.

In the pursuit of identifying the most effective presentation method for movie recommendations, Nanou, Lekakos and Fouskas (2010) conducted an in-depth analysis of past research approaches and popular movie recommendation systems. By comparing multiple methods of presentation, the study was able to identify the most effective approach based on user satisfaction and persuasion, and demonstrated a positive interrelationship between the two across all empirical conditions. The study also highlighted the significance of emotional and aesthetic factors in the decision-making process of users.

Given that the way choices are presented to individuals can have a profound impact on their decision-making, this study delves into the choice architecture behind human-computer interactions. Through this examination, the study aims to enhance the relevance of diversity items for users.

2.2 A Nudge Theory Approach to Decision-Making

Recommender systems have a crucial function beyond just offering advice to users. To fully grasp the potential impact of these systems, it is vital to perceive them not only as social agents but also as persuasive agents (YOO; GRETZEL, 2011). While satisfaction is a commonly measured parameter in RSs, along with other aspects such as trust and acceptance of recommendation technologies, persuasiveness is often overlooked (NANOU; LEKAKOS; FOUSKAS, 2010).

Persuasive Recommender Systems (PRS) is a subclass of RS that alter the user's behavior or attitude through interaction with the system (YOO; GRETZEL; ZANKER, 2013). Empirical studies have proven that persuasion theories developed for human-human interactions can be implemented in technological contexts, leading to more persuasive interactions. By incorporating these theories, it is possible to understand the role of PSR and its interplay with the user.

As a result of human-computer interaction, RSs persuade users and influence human behavior (YOO; GRETZEL; ZANKER, 2013). Therefore, the user's decision-making process may not always be rational, as heuristics and biases can play a role (TVERSKY; KAHNEMAN, 1974). This indicates that cognitive limitations can impact rational decision making (LEX *et al.*, 2021).

Prior research has highlighted the importance of recommendation content and format on user behavior and satisfaction with the system. The interface design and the information format presented are critical factors during the development of persuasive systems (YOO; GRETZEL; ZANKER, 2013). In this study, the nudge theory was applied to design the interface, implement persuasion techniques, and develop the recommendation application.

RSs can be viewed as nudges from a behavioral economics standpoint since they involve individual decision-making. The concept of nudges is based on the idea that psychological phenomena influence human decisions. A choice architect can leverage these phenomena to create a choice environment that subtly guides people towards making what they perceive to be the best choice (JESSE; JANNACH, 2021). In digital contexts, this concept is known as digital nudging, which utilizes user interfaces to influence people's digital decision-making (WEINMANN; SCHNEIDER; BROCKE, 2016).

As the choice architect in a digital environment, it is the responsibility of an RS to understand how even the slightest details can influence people's decisions and remember that "everything matters." Thaler and Sunstein (2009) define nudges as a way of altering people's behavior, and these nudges are classified as libertarian paternalism. Based on the idea that people should be able to make their own decisions, nudges are designed to help individuals make more informed and beneficial decisions without being intrusive or limiting.

In the context of developing recommender systems, it is important to consider how nudges can guide user behavior while still allowing them to make their own choices. According

Nudging Mechanisms			
1) Decision Information	2) Decision Structure	3) Decision Assistance	4) Social Decision Appeal
Translate information	Change range or composition	Provide reminders	Increase reputation of messenger
Increase salience of information	Change choice defaults	Facilitate commitment	Provide social reference point
Make information visible	Change option consequences		Instigate empathy
Change phrasing of information	Change option related effort		

Table 1 – Nudging Mechanisms. Source: (JESSE; JANNACH, 2021)

to Meske and Potthoff (2017), design, information, and interaction elements can all be used to create effective nudges in digital environments.

To further explore this concept, a taxonomy of nudge mechanisms specifically for recommender systems was developed by Jesse and Jannach (2021). This taxonomy consists of four primary categories and thirteen subcategories, which are detailed in Table 1. Next, we will discuss the four primary categories.

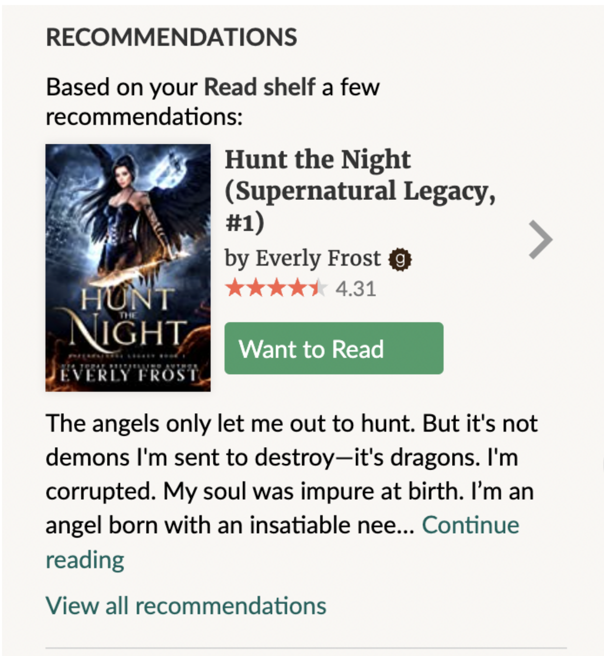

Figure 2 – Example of of Decision Information Category

Figure 3 – Example of Decision Structure Category

Decision Information category aims to modify the presentation of information without altering the choices available. An example of this category can be seen in Figure 2, taken from the Goodreads website, where a book recommendation is accompanied by information about why the book was recommended.

Decision Structure category involves modifying the decision-making process to influence decision makers. Nudging mechanisms that focus on modifying the arrangement of options fall under this category. An example of this category can be seen in Figure 3, sourced from the Amazon website, illustrates a list of book recommendations filtered by the most popular one. The structure of the list has been modified to include two sponsored books. Despite not being the most popular, both books belong to the LGBTQIA+ category and are relevant to the list.

Decision Assistance is focused on supporting decision makers to achieve their goals. For instance, Figure 4, taken from the Goodreads website, provides an example of this category in the form of a notification urging the user to set a goal for the number of books they want to read in 2023.

Social Decision category entails redesigning the information presented to users with a focus on emotional and social implications. This hypothesis is based on the assumption that users are influenced by what others have done in similar circumstances. An example of this category can be seen in Figure 5.

The taxonomy has identified a total of 87 nudging mechanisms, as per Jesse and Jannach (2021). However, the survey conducted by the authors revealed that 69 of these mechanisms have yet to be thoroughly investigated.

Moreover, the survey findings suggest that nudging mechanisms can be effectively employed in recommender systems to address biases. Many studies have demonstrated the positive impact of nudging mechanisms on users' behavior.

2.3 Final Considerations

Recommendation systems are known to face various biases that can significantly impact their effectiveness. Identifying and addressing these issues has become a key challenge for researchers striving to improve recommendation systems.

However, some approaches described in this chapter suggest that exposing users to a high level of accuracy and limiting interactions can create filter bubbles and unfairness, posing significant challenges to the state of the art. While filter bubbles are a practical concern, failure to recognize and address them can have serious consequences.

Although nudge mechanisms are already being employed in RSs, there are still gaps in how we can utilize these mechanisms to encourage users to engage more with diverse items.

In light of these challenges, this thesis project explores the use of nudge mechanisms to address filter bubbles and promote diversity exposure. The next chapter examines relevant literature on biases and filter bubbles, as well as research that employs nudge mechanisms in recommender systems.

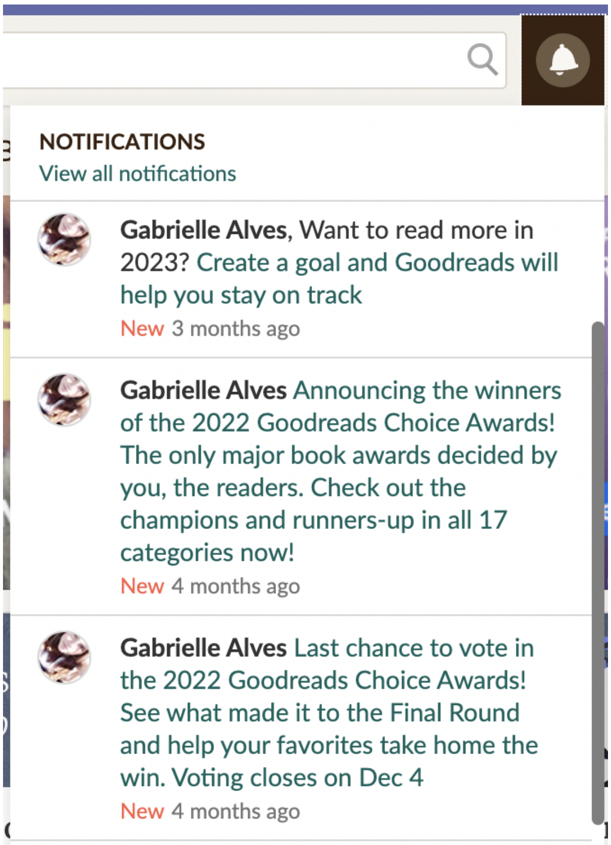

Figure 4 – Example of Decision Assistance Category

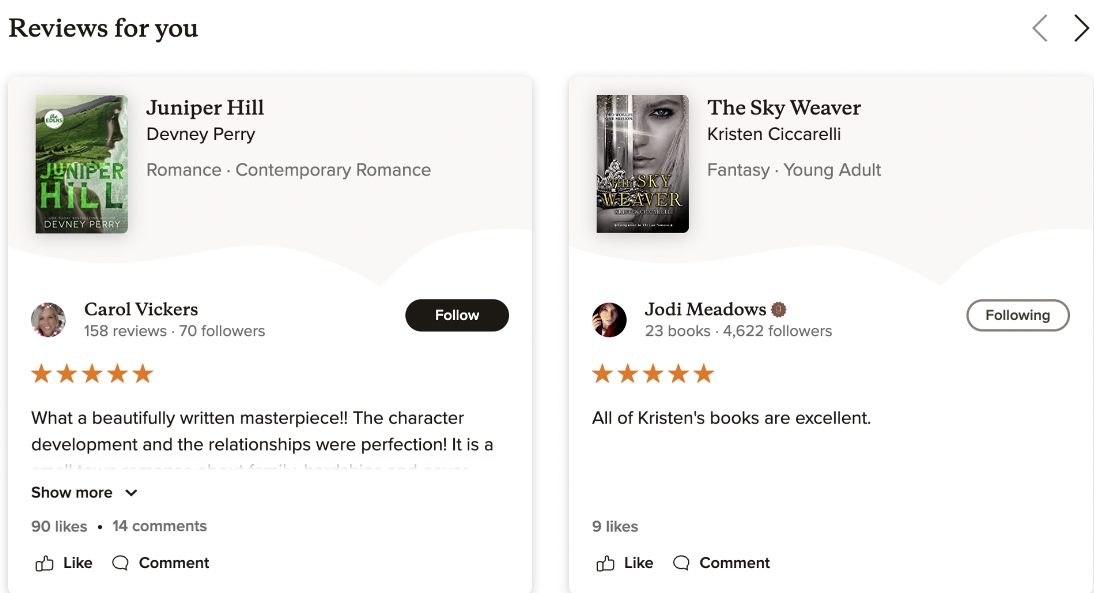

Figure 5 – Example of Social Decision Category

CHAPTER

3

RELATED WORK

In this section, we present a summary of related literature aligned with the specific objectives of this work. We aim to identify open questions and contribute to state-of-the-art innovation by addressing them.

To conduct this review, we utilized a virtual research environment that curates academic literature and metadata. We conducted keyword searches using terms such as "System Recommendation," "Diversity Exposure," "Filter Bubble," "Bias and Unfairness," "User-Centric Interface," and their synonyms to retrieve relevant articles. By filtering the results by time period, we identified recent works that are pertinent to our study.

Our review will also explore strategies that utilize nudge mechanisms in recommendation systems, as proposed by Jesse and Jannach (2021).

3.1 Exposure Diversity as a Design Principle

Recommender systems aim to help users discover relevant items that were previously unknown to them. However, the extent to which different algorithms are able to surface content outside a user's past preference profile varies. Content-based filtering systems (LOPS; GEMMIS; SEMERARO, 2011), for instance, recommend items that are similar to the user's past profile. In contrast, collaborative filtering has more potential for discovering items outside a user's familiar taste profile by relying on preference patterns in the user community. This approach is the focus of our present work.

Previous research found that recommender systems can indirectly guide users to categories in which they had not previously made purchases. For instance, (DIAS *et al.*, 2008) and (LAWRENCE *et al.*, 2001) reported such findings in online supermarkets, while (KAMEHKHOSH; BONNIN; JANNACH, 2019) reported similar observations in the music domain.

Finally, reinforcement learning-based approaches use explore-exploit schemes to gather more information about users' preferences. Such recommender systems may intentionally recommend items for which they are uncertain about the suitability for the users (MCINERNEY *et al.*, 2018).

Regardless of the algorithm used, it has been widely recognized that simply recommending items with high predicted relevance may not be sufficient (MCNEE; RIEDL; KONSTAN, 2006). Therefore, various approaches have been proposed in recent years to increase the diversity, novelty, and serendipity of recommendations for users.

In a recent large-scale study, Chen *et al.* (2019) investigated the perception of novel, serendipitous, and diverse recommendations among over 3,000 customers of a Chinese mobile e-commerce platform. The study found that serendipity significantly influenced user satisfaction and purchase intent, and that it had a more direct effect than diversity or novelty. Our work is similar to theirs, in that we study human perceptions involving users of an online service, but we focus on the effects of nudging users towards off-profile items.

We note that increasing serendipity, diversity, and novelty can potentially help users discover relevant, novel content. However, the extent to which diversity-aware or novelty-aware recommender systems support this goal depends on how diversity and novelty are operationalized, such as the specific evaluation metric or optimization goal. Including items of different genres in the recommendations, for example, may increase metrics such as intra-list similarity (ZIEGLER *et al.*, 2005), but it does not guarantee that these genres are new to users (off-profile).

In some works, novelty is defined in terms of an item's general popularity, but such an approach does not guarantee the discovery of off-profile content. Determining novelty in terms of an item's distance to a user's past profile, as in (CASTELLS; VARGAS; WANG, 2011), could lead to the discovery of off-profile content. Similarly, serendipity or unexpectedness is also difficult to operationalize in computational metrics.

Several research works have specifically focused on the exploration and discovery aspect of recommender systems. For instance, in (KAPOOR *et al.*, 2015), the authors examine how users' openness to explore new musical content changes over time. By analyzing data, they confirm the existence of such dynamics and develop a method to predict variable novelty preferences. Unlike our study, the authors consider novelty at the item level, using listening history to identify novel items. In contrast, our work investigates if users can be encouraged to explore off-profile items through a user study. Furthermore, we approach the research question differently from (KAPOOR *et al.*, 2015), as we explore the users' behavior through a user study.

The domain of music also tackles discovery questions, as highlighted in a user study conducted by Ludewig *et al.* (2021). The authors evaluated multiple recommendation algorithms for creating dynamic playlists and found that Spotify's recommendations outperformed all others in helping users discover previously unknown tracks that they enjoyed. Interestingly, although Spotify's offline accuracy was low, users still reported a high intention to use the system again and recommend it to others. This suggests that supporting discovery in this domain is crucial and valuable.

The significance of these topics for the practical application's business success is also emphasized in recent research by Spotify's scholars, where the explicit goal is to nudge users' consumption towards less popular and diverse content that deviates from their past preferences (HANSEN *et al.*, 2021). The authors assess the distance between a given track and a user profile by comparing track and profile embeddings. They observe a trade-off between diversity, which in their study means both recommending less popular and more unfamiliar items, and user satisfaction. Their computational analyses aim to identify suitable algorithmic approaches to balance this trade-off. In contrast to our work, the authors of (HANSEN *et al.*, 2021) rely solely on offline experiments. They use historical listening logs as a proxy for user satisfaction and consider a user dissatisfied if they skip a recommended track. Nevertheless, similar to our work, they also explore a diversification strategy where highly predicted relevant tracks are interleaved with highly diverse tracks.

Heitz *et al.* (2022) conducted a user study that highlighted the potential societal impact of diversifying content in certain domains of recommender systems beyond users' typical preferences. For instance, in their study, the authors developed a mobile news reader app with a customizable algorithm that diversified news articles based on political orientation. The authors evaluated different diversification strategies through a field test and found that diverse news recommendations did not harm the user experience. In fact, their results suggested that recommendations enhance tolerance for opposing views and may have a depolarizing capacity for democratic societies. Similarly, our present work aims to guide users outside of their filter bubble in mobile app recommendations, but we actively try to nudge users towards off-profile items. Unlike Heitz *et al.* (2022), our work focuses on the movie domain and uses a user study to investigate whether nudging users towards unfamiliar items affects their experience.

3.2 Nudging Towards Diversity

In 2008, Thaler and Sunstein (2009) popularized the concept of nudge theory, which sees nudging as a way to positively influence people's behavior without coercion. Nudging can be applied in various real-world scenarios, such as the placement of healthy food in a cafeteria to encourage healthier eating habits. Similarly, nudging can also be applied in the digital realm, often referred to as digital nudging (CARABAN *et al.*, 2019; SCHNEIDER; WEINMANN; BROCKE, 2018). A typical example of digital nudging is the pre-selection of a default option, such as when users are given a choice between different subscription plans for a service.

Caraban *et al.* (2019) provided a comprehensive review of digital nudging techniques in human-computer interaction research, where they identified 23 ways of nudging that were organized into six groups. Interestingly, the authors also acknowledged that some of the identified nudges could be used in a non-benevolent way to deceive users, which stands in contrast to the original conception of nudges, which were commonly designed for benevolent purposes. Some of the identified nudges could therefore rather be viewed as covert *persuasion*. Persuasion techniques, which are related to nudging, are often also based on certain psychological phenomena that can be used to influence human behavior (FOGG, 2002). Mols *et al.* (2015) also see nudges as related yet different concepts, where nudges are effective when targeting behavior change in a particular setting, while persuasion targets underlying beliefs and preferences and should lead to a more sustained behavior change.

In their comprehensive review, Jesse and Jannach (2021) explored the topic of nudging in combination with recommender systems. Through an exhaustive analysis of the literature, the authors identified 87 nudging mechanisms and discussed the various psychological phenomena that explain why nudges are effective. They also argued that recommender systems themselves can be seen as a tool for nudging users, as they commonly use digital nudging techniques such as increasing item salience, leveraging positioning effects, and simplifying access to certain items. However, the literature on this topic is not always consistent, and terms such as "hidden persuaders" may be used to describe recommendations that are not targeting underlying beliefs. For a more in-depth discussion on the intersection of persuasion and recommender systems, interested readers may refer to Yoo, Gretzel and Zanker (2013)'s work.

Our study explores the combination of recommendations and nudging by applying nudges to specific items in a recommendation list. Previous research has examined the use of nudging in search and recommender systems, particularly in the food (recipe) domain. For example, Elsweiler, Trattner and Harvey (2017) demonstrated through a user study that selecting suitable food images can nudge users towards choosing healthier recipes. Similarly, Starke, Willemsen and Trattner (2021) found that participants tended to select healthier options when they had visually attractive images attached. Jesse, Jannach and Gula (2021) explored the effectiveness of different types of nudges in the food domain and found that a hybrid nudge, which combined setting a default and providing a social cue, was the most effective in influencing food choices without negatively impacting user satisfaction.

The work presented in Starke, Willemsen and Snijders (2020) examined the potential of nudging in the energy domain by implementing a "social norm" nudge in a recommender system designed to suggest energy-saving measures. While the results were mixed, and not all hypotheses were confirmed, the study made some critical observations. For instance, the perceived feasibility of a measure, i.e., how difficult it would be to implement, was found to mediate the effect of the nudge.

Liang and Willemsen (2021) investigated the effectiveness of digital nudging in encouraging users to explore music genres that are outside their current preferences. The authors conducted an exploratory study and found that users are more likely to select distant genres if they appear at the top of the list due to order effects and position bias. In contrast to their work, our study differs in that we interspersed off-profile items throughout the list in both the control and treatment groups. Furthermore, we applied additional nudges in the treatment group to encourage users to engage with off-profile items.

In a recent study by Vermeulen (2022), the potential role of recommender systems in achieving public policy goals is discussed. Article 10 of the European Convention on Human Rights requires governments to ensure access to diverse news sources and information via audiovisual media for citizens, which has not yet been established for digital media. Vermeulen proposes that recommender systems and digital nudging can be used to promote online access diversity and reduce filter bubbles and echo chambers. These effects are often the result of confirmation bias, where online users tend to consume content that aligns with their prior beliefs. In a similar vein, Rieger *et al.* (2021) explored how digital nudging could mitigate confirmation bias and aid users in making informed decisions. Our study also aims to help users discover relevant content outside their preferences, but in a distinct domain.

In previous studies, there have been instances where recommender systems were combined with nudging or persuasive techniques. For instance, these techniques were used in recommending environment-friendly routes (Efthimios Bothos; Dimitris Apostolou; Gregoris Mentzas, 2016), news (GENA *et al.*, 2019), and social media followers (VERMA *et al.*, 2018). For more information on this topic, see also (JESSE; JANNACH, 2021). However, to the best of our knowledge, there have been no previous attempts to apply nudging in the context of book recommendation, which is the focus of our current research.

3.3 Impact of Diversity on the User's Experience

Recommending items of various subjects can broaden the users' interests beyond popular items. However, this approach also raises the risk of suggesting items that are completely new or unfamiliar to the user, resulting in reduced accuracy and affecting the user's experience and satisfaction with the system (WERNECK *et al.*, 2021).

Recommender systems (RSs) can be classified as social actors as users interact with the system socially, and should be studied from both technical and social perspectives. Hence, the effectiveness of RSs extends beyond the quality of the algorithm. Ensuring user acceptance of RSs poses several challenges, including enhancing user satisfaction with the system. Research suggests that the content and format of recommendations significantly impact user behavior and satisfaction with the RSs (YOO; GRETZEL; ZANKER, 2013).

The selection, organization, and presentation of information play a vital role in the effectiveness or failure of recommendations. To explore the relationship between users' perception and the system's trustworthiness, Chen. and Pu. (2005) developed a trust model and conducted a survey. This is important because trust is a significant factor that influences a user's decision-making process.

In their investigation of customer satisfaction with a deep learning-based recommender system, Kim, Choi and Li (2021) analyzed various recommendation approaches in terms of their accuracy and diversity. The study revealed that both accuracy and diversity have a positive impact on customer satisfaction. In contrast, traditional recommender systems seem to positively affect customer satisfaction only through their accuracy.

Based on the findings of Nguyen *et al.* (2018), rating-based recommender systems often fall short in providing the level of diversity, popularity, and serendipity desired by their users. Moreover, users with varying personalities have distinct preferences for these factors. The study suggests that incorporating users' personality traits into the recommendation generation process can enhance user satisfaction.

Our study seeks to expand on the previous research by investigating the factors that impact users' decisions, perceptions, and trust when engaging with diverse content. Additionally, we aim to explore the potential impact of digital nudging techniques on several measures, including satisfaction, trust, and effectiveness.

3.4 Final Considerations

This chapter delves into various approaches to promote diversity exposure in recommendation systems. While accuracy algorithms are important, other factors such as user satisfaction and trust should also be considered in building a sustainable system. However, current approaches tend to focus more on developing algorithms and metrics, without fully exploring the impact of exposing users to unfamiliar or novel items on evaluation metrics like satisfaction and trust.

To address this, our proposal takes a human-centric approach to diversity exposure. We aim to encourage users to interact with relevant but unfamiliar items through the use of digital nudging and user experience evaluation. The following chapter outlines our study design in detail.

CHAPTER

4

USER STUDY

This chapter presents the user study conducted as part of the master's research and is structured as follows. Firstly, in Section 4.1, the research gaps and objectives are thoroughly presented, highlighting the significance and relevance of the study's objectives to the existing literature. Secondly, Section 4.2 provides an overview of the project's methodology, including experiment design, data collection, and analysis techniques.

4.1 Study Overview

The focus of our research is to investigate the impact of digital nudging techniques on actively directing users towards off-profile recommendations. Specifically, we aim to study the effectiveness of different nudging mechanisms in encouraging users to explore and consider these recommendations, and how these nudges affect their perception of the quality of recommendations and their future behavioral intentions.

Our study is centered around the domain of book recommendations and utilizes a real-world Brazilian social book recommendation site. The platform's primary goal is to increase literacy levels within society, making it an ideal target domain for our work. Through this research, we aim to provide insights into effective strategies for promoting diversity in users' reading interests and engaging them in reading, ultimately contributing towards the enhancement of literacy levels.

Our study employed a between-subjects design, where participants were randomly assigned to either the treatment or control group. Both groups used a mobile book recommendation app developed for the study and received recommendations containing books from their preferred genres as well as off-profile books. However, the treatment group had an enhanced user interface that included different types of digital nudges attached to each off-profile book. Participants in both groups were asked to select books to add to their reading lists, and we analyzed both their objective behavior and subjective statements regarding the quality of the recommendations and choice process. This approach allowed us to evaluate the effectiveness of the digital nudges in encouraging users to consider off-profile recommendations and to gain insights into the participants' perceptions of the recommendation process.

It is important to note that we utilized our own app for the study for two primary reasons. Firstly, we prioritized the user experience, and a mobile app was the most convenient way to serve the target users. Secondly, while we aimed to provide a realistic user experience, we wanted the participants to be aware that they were interacting with the app for research purposes.

4.2 Study Design

4.2.1 Experiment Flow

Experiment Flow is illustrated in Figure 6. We recruited study participants via social media who downloaded a mobile app designed for the study. Upon opening the app, participants were informed about the study's purpose and tasks. The study's cover story informed participants that the goal was to develop a better book recommendation system for avid readers and those seeking to establish a reading habit. Participants were able to create a profile and specify their favorite genres in the app. The app then provided book recommendations from both the participant's preferred genres and off-profile genres. The participants were encouraged to select between three to five books to add to their reading list, but were not required to select a specific number of books. We randomly assigned the participants to either the treatment or control group.

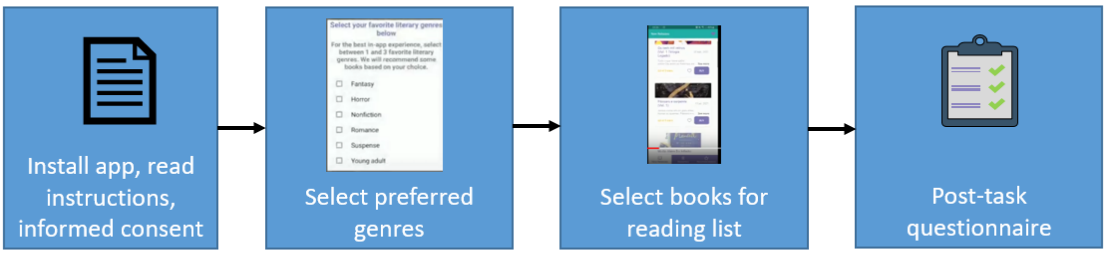

Figure 6 – Experiment flow.

Upon providing informed consent, participants were asked to select their preferred book genres from a list of six. They were then presented with a list of recommended books, with every second book belonging to a non-preferred genre. Participants were free to add as many books as they desired to their reading list. Following their selection, participants received an email requesting that they complete a post-task questionnaire. The questionnaire included questions regarding their perception of the recommendations, feedback about the recommendations, and demographic questions. To encourage participants to complete the survey, those who did were automatically entered into a lottery draw, with prizes consisting of bookmarks and books.

4.2.2 Preference Elicitation & Recommendation List Design

For the preference elicitation phase of the study, participants were presented with a list of six genres: *Fantasy*, *Horror*, *Nonfiction*, *Romance*, *Suspense*, and *Young Adult*. These categories were chosen based on their popularity in the books section of Amazon.com.br at the time of the study. Our aim in selecting these popular categories was to increase the likelihood that participants would express a preference for at least some of them.

The recommended books for each category were carefully selected based on multiple factors, including reader ratings and new releases from independent publishers and debut authors. This selection process was aimed at ensuring that the recommender could facilitate the discovery of new books and was not restricted to only bestsellers. For further details, please refer to the Appendix, which provides the complete list of recommended books.

Participants in both the treatment and control groups were presented with a total of 24 book recommendations, half of which were selected from the preferred genre category they had previously indicated. The order of the books was randomly generated and kept consistent across all users. The remaining half of the book recommendations were chosen from non-preferred categories.

To ensure a consistent user experience, we maintained a static interleaving rule and number of books on the recommended list for all participants. However, the specific recommended books were tailored to the genres selected by each user. For instance, if a user selected four genres, we interleaved twelve random books from the preferred genres with twelve random books from the two non-preferred genres.

Participants who selected all six categories as their preferred genres were excluded from the analysis, as they were only presented with books from their preferred genres and thus did not experience the full range of recommended books.

To avoid order effects and presentation biases, we adopted an interleaving approach inspired by previous studies such as (JOACHIMS, 2002) that addressed the issue of click-through analysis in information retrieval settings. Previous research has shown that items displayed at the top of a list are more likely to be seen and clicked by users, which may not accurately reflect their relevance (BAEZA-YATES; RIBEIRO-NETO, 2011) (OOSTERHUIS, 2020). As a result, comparing two ranked lists can be challenging. To mitigate this issue and obtain more unbiased feedback from users, we used an interleaving method where preferred items are interspersed with off-profile items. This means that every second item in the list is not from a user's preferred genre.

The interleaving approach adopted in this study is illustrated in Figure 7, where items of preferred genres are interleaved with off-profile items. It is important to note that this interleaving was implemented for both the control and treatment groups. The only difference between the two groups is that a digital nudge was attached to each off-profile item in the treatment group.

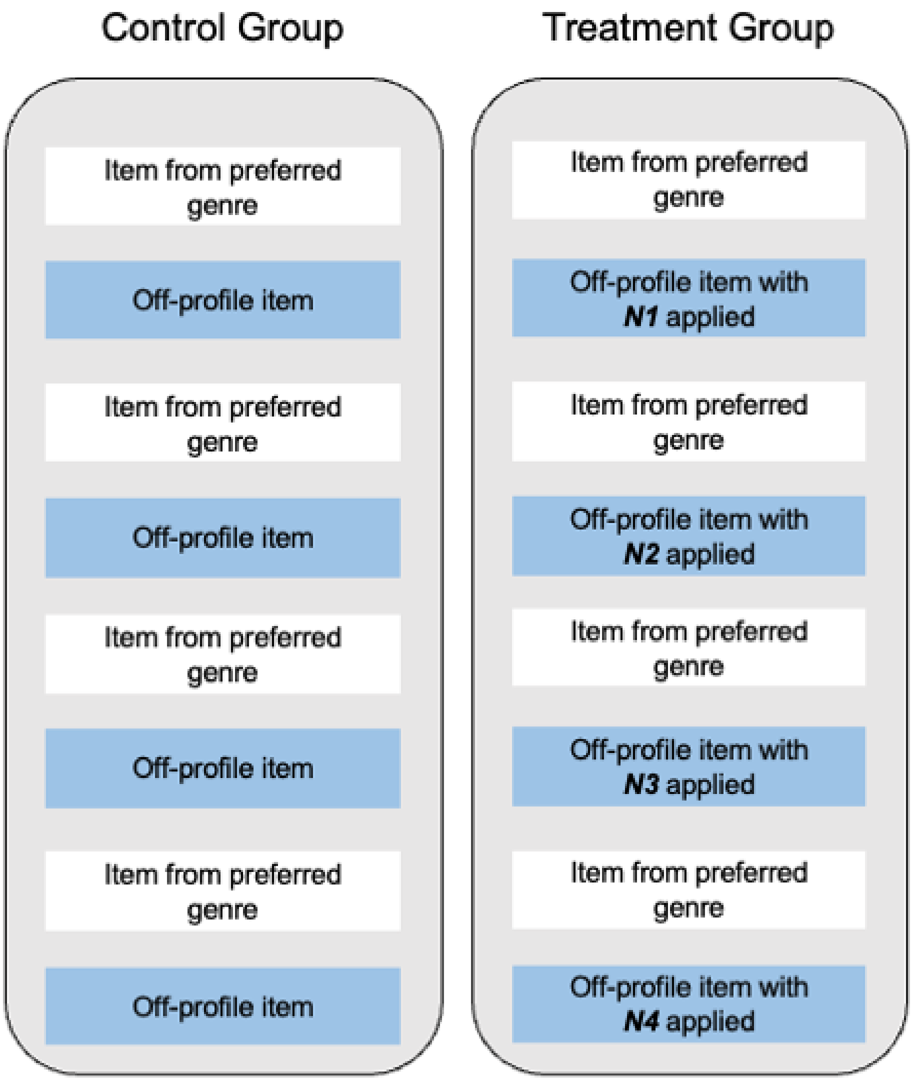

Figure 7 – General Structure of Interleaved Recommendation Lists (only showing the first four of 24 items).

4.2.3 Selected Nudges

We examined *four* types of digital nudges based on existing literature, which are commonly used in practical applications. As our experiments were conducted in the context of a social book recommendation site, we focused on social nudges (JESSE; JANNACH, 2021), which are rooted in social psychology theories like social comparison and conformity (GENA *et al.*, 2019). According to these theories, individuals who are uncertain about how to behave in a given situation seek information on what others have done in similar situations. This information then influences their attitudes and behavior. The details of each nudge are outlined below.

- *Nudge N1 "Hybrid: Following the herd and Increase Salience*": One example of a nudge used in the literature (THALER; SUNSTEIN, 2009; MIRSCH; LEHRER; JUNG, 2017) is to display a message to users indicating that others have favored or selected a specific item. The underlying hypothesis is that users are more likely to follow the trend of the majority since they do not want to stand out from the crowd. Our implementation of this nudge visually highlighted the message, indicating that many other people had added the book to their reading list, e.g., *"Favorited 36 times today"*[1]. We chose this nudge since it is commonly used in practical applications, such as on flight reservation websites.

- *Nudge **N2** "Hybrid: Increase salience of attribute and Argumentum-Ad Populum"*: This nudge was designed to enhance the visibility of the targeted item and attract more attention from the users. Additionally, we incorporated a message based on the hypothesis that individuals may be more likely to accept an opinion if the majority shares it. For example, we included a statement such as *"90% of users liked this book"*[2]. We chose this hybrid nudge because recent research in the food industry demonstrated its high effectiveness (JESSE; JANNACH; GULA, 2021). Displaying the number of *likes* is commonly used in various practical applications, including social media platforms.

- *Nudge **N3** "Hybrid: Increase salience of attribute and Messenger Effect"*: This nudge is a modification of **N2**. In this case, rather than emphasizing the number of users who liked a book, we presented messages such as "Netflix will adapt this book" or "Bestseller of the week on Amazon." Such labels are prevalent in the offline world as well, such as "New York Times Bestseller."

- *Nudge **N4** "Social reference point"*: This nudge pertains to the influence of opinion leaders on a book's evaluation. The underlying hypothesis is that users often rely on the opinions of famous and influential individuals. To account for participants' diverse backgrounds and demographics, we compiled a list of opinion leaders, including George R R Martin, Neil Gaiman, Obama, Stephen King, Lupita Nyong'o, Reese Witherspoon, Emma Watson, Dakota Fanning, Mark Zuckerberg, and Bill Gates. An example message in this nudge was *"Recommended by Stephen King"*. It is worth mentioning that leveraging endorsement statements by opinion leaders or celebrities is a common tactic in traditional marketing.

The order of the four nudges was randomized before the experiment and remained unchanged throughout the study. As each recommendation list comprised 12 items with nudges in the treatment group, the static order was consistently applied after the fourth and eighth off-profile item. It is worth mentioning that all nudge statements concerning the ratings and opinion leaders were authentic, except for nudge N1, where the number of recent favorites was fabricated.

We would like to emphasize that our choice of "Social Decision Appeal" category nudges from Jesse and Jannach (2021) was motivated by our application setting. We enhanced

[1] It is important to note that the app was in Portuguese; the provided examples are translated to English.
[2] To increase realism, we slightly varied the numbers used in the message. For further details, see (JESSE; JANNACH; GULA, 2021).

the salience of information in three cases using a hybrid nudging approach, which has demonstrated its effectiveness in previous studies (JESSE; JANNACH; GULA, 2021; RENAUD; ZIMMERMANN, 2019). While there are several other types of nudges in the literature that could be applied, we decided to focus solely on one category of nudges in our current work. This is because we did not aim to compare the effectiveness of different nudge types, as has been done in previous research such as (BERGER; MÜLLER; NÜSKE, 2020; JESSE; JANNACH, 2021).

We acknowledge that the categorization of nudges in the literature, such as the "Social Decision Appeal" category in Jesse and Jannach (2021)'s work, can sometimes lack clear boundaries. Additionally, a single type of nudge can be linked to multiple psychological phenomena, as Jesse and Jannach (2021) have also noted. In our study, some nudges display similarities, such as Nudge N3 and N4, which both involve an authority figure in some way. Nevertheless, all four nudges utilized in our research incorporate a social aspect.

4.2.4 Objective Measurements of User Behavior

In order to analyze the impact of the nudges on users' decision-making, we collected various data points. Firstly, we recorded the number and specific titles of the books that the participants added to their reading lists. Additionally, we documented the positions of these books, as well as whether they were preferred choices or off-profile items. To gain deeper insights into participants' behavior, we also tracked the number of books they viewed in detail (by clicking on a "show more" button) and how far they scrolled down the list. Finally, we recorded the amount of time participants took to make their selection.

4.2.5 Assessing Subjective User Perceptions

To address our second research question regarding the impact of the nudges on users' perceptions, we included relevant questions in the post-task questionnaire. Our questionnaire is based on the Recommender systems' Quality of user experience (ResQue) framework proposed by Pu, Chen and Hu (2011), which is a widely used and validated instrument for evaluating user perceptions of recommender systems. The ResQue framework connects various factors, including user-perceived quality factors, user beliefs, user attitudes, and behavioral intentions. At the first level, user-perceived quality factors, such as accuracy and diversity of recommendations, are linked to user beliefs at the second level, such as perceived usefulness, transparency, and ease-of-use. These beliefs may impact user attitudes at the third level, such as satisfaction and trust, which in turn may influence behavioral intentions at the fourth level, such as intended purchases or use. An overview of the ResQue Framework can be seen in Figure 8.

We utilized the ResQue framework's validated questionnaire as a foundation for our study, although we made adaptations to suit our mobile setting with small screen sizes and high user expectations for app usability. As such, we only included a subset of the original

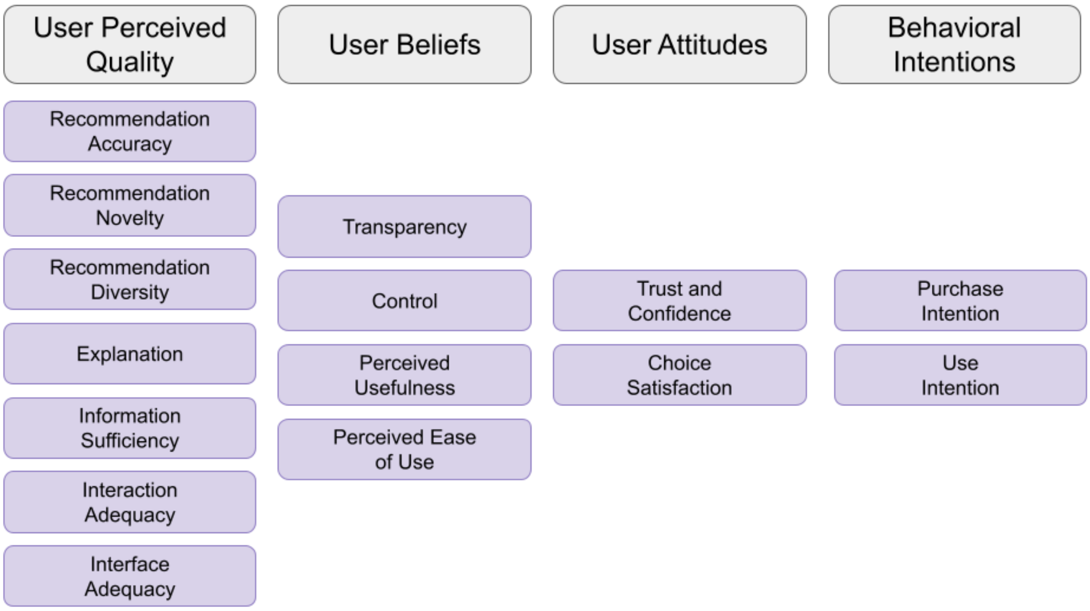

Figure 8 – Overview of the ResQue framework for User-Centric Evaluation of Recommender Systems (see (PU; CHEN; HU, 2011))

questionnaire, with specific questions for each dimension presented in Table 2[3]. Participants provided responses on a five-point Likert scale, with 1 indicating complete disagreement and 5 indicating complete agreement. To avoid overwhelming participants in the mobile app, we measured the constructs shown in the table with only one questionnaire item. It is important to note that using only one item per construct may represent a research limitation. However, we believe that the potential risks, such as misinterpretation of the questions, are minimal as we used questions from the ResQue framework's validated survey.

Furthermore, we incorporated a set of questionnaire items aimed at gathering information about the participants' demographic characteristics and their overall interest in reading books, in addition to the items focusing on their perceptions.

4.2.6 Technical Implementation

The main screen of the application is displayed in Figure 9, which appears after the participants have provided their preferred genres. The application consists of two tabs, with the first tab displaying a list of recommendations. The screen capture illustrates the application of nudge N4 to one of the books recommended, which is labeled as "Recommended by Stephen King".

Participants were able to access more information about the book by clicking on "show more." Upon doing so, they would be directed to a details page where they could view the full book cover, title, author, publisher, price, and a brief synopsis. An example of a book details

[3] The questions were translated into Portuguese

ID	Construct	Question
Q1	Accuracy	The items recommended to me matched my interests.
Q2	Attractiveness	The app provided some really good recommendations.
Q3	Diversity	The items recommended to me are diverse.
Q4	Novelty	The recommendations helped me discover new books which I find interesting.
Q5	Serendipity	I was surprised by some of the recommendations.
Q6	Information Sufficiency	The amount of information available was sufficient to make the decisions.
Q7	Transparency	I understood why the items were recommended to me.
Q8	Usefulness	Finding books to put on my reading list was easy.
Q9	Choice Satisfaction	My overall satisfaction with my selection of books on my reading list is high
Q10	General Satisfaction	My overall satisfaction with the app is high.
Q11	Intention to use	I like the app, and I would use it again in the future.
Q12	Intention to "purchase"	I am eager to read the book(s) I have put on the reading list.

Table 2 – Question Items in the Post-Task Questionnaire

page is displayed in Figure 10.

To add an item to the reading list, participants simply had to click on the heart symbol. The second tab of the application allowed them to view their currently favorited items, i.e., the ones added to their reading list.

To conduct the experiments, we developed an Android application in the Kotlin language, with a focus on the Android platform for its wide user base and easy distribution through the Google Play app store. The application followed the Clean Architecture (MARTIN,) and Model-View-View Model (MVVM) design pattern, resulting in a maintainable, testable, and readable architecture. To simplify list updates, the back-end used a basic online document. The same software built both the Treatment and Control versions of the application, with the Treatment group displaying the relevant nudges through an enabled switch. All user interactions were recorded using the Mixpanel tool [4], streamlining the logging and analysis of application log events.

[4] <https://mixpanel.com>

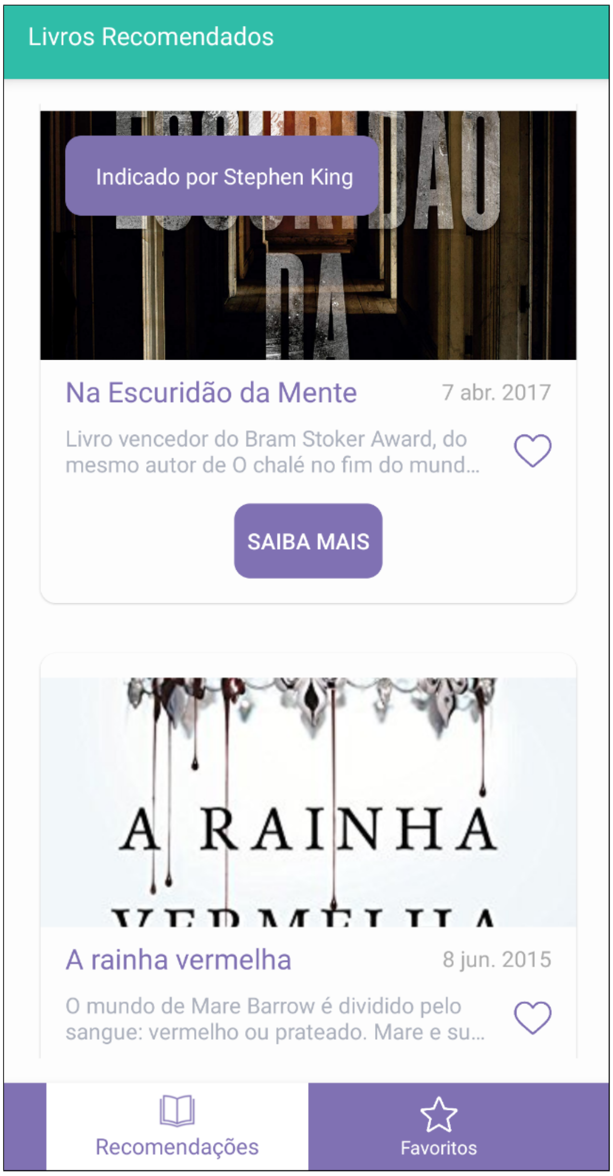

Figure 9 – Screen Capture taken from the Treatment Group, with nudge N4 applied (in Portuguese).

4.3 Final Considerations

This chapter provides an overview of the research objectives, design, and methodology used to investigate the impact of digital nudging techniques on promoting diversity in users' reading interests and engagement in reading.

In the next chapter, we will present the results of the study and discuss their implications for the use of digital nudging in promoting off-profile recommendations, and we examine how the different types of digital nudges affected participants' behavior and perceptions of the recommendation process.

Figure 10 – Screen Capture of a Book Details Page.

CHAPTER

5

RESULTS & DISCUSSION

This section summarizes the major findings of our study. We begin by providing an overview of our study's execution and participant recruitment process. Next, we delve into our investigation of the impact of nudges on both user behavior and perceptions. By exploring these key areas, we have gained valuable insights into the effects of nudges and their potential implications, contributing to a deeper understanding of this important topic.

5.1 Study Execution & Participants

The study was conducted over a period of approximately four weeks, between 11 February 2022 and 08 March 2022. To recruit participants for our study, we reached out to users of the popular Brazilian social book recommendation site "Livros & Citações"[1] through various channels, such as social media platforms like Instagram. Interested participants were required to download and install the app from the Google Play app store in order to participate in the study. Our recruitment strategy successfully resulted in a diverse group of participants who were able to provide valuable insights into the effects of nudges on user behavior and perceptions.

Our study included a total of 1,064 participants, all of whom added at least one recommended book to their reading list using the app. Of these participants, 520 were assigned to the treatment group, while 544 were placed in the control group. A total of 762 subjects successfully completed the post-task questionnaire, with 367 in the treatment group and 395 in the control group. The majority of our participants, over 90 %, identified as female. When considering age demographics, we found that more than 50 % of subjects were between the ages of 18 and 25, which is consistent with the average age range of users of the book recommendation site. Less than 10 % of participants were over 40 years old. Additionally, the majority of participants could be classified as engaged book readers, with over 65 % stating they read between 6 to 50 books per year, and over 15 % reading even more than 50 books per year. Further details regarding

[1] <https://www.livrosecitacoes.com>

participant demographics, including gender and reading habits, can be found in Table 3. It is worth noting that the differences in these characteristics between the control and treatment groups were not statistically significant (as determined by a Student's t-test with α=0.05), ensuring that any observed effects of nudges on user behavior and perceptions were not biased by differences in these demographics.

Our study aimed to achieve highly reliable results by using a sufficiently large sample size for a robust analysis. This is especially important when utilizing Structural Equation Modeling (SEM) techniques, as we did in our analysis. These techniques typically require a larger number of samples to yield accurate results, as noted by Christopher Westland (2010) in discussions on the lower bounds of sample sizes.

Variable	Control	Treatment
Gender		
Female	95,19 %	93,73 %
Male	4,05 %	5,99 %
Other	0,76 %	0,27 %
Reading habits		
1-5 books per year	8,10 %	7,36 %
6-20 books per year	33,16 %	38,42 %
21-50 books per year	34,94 %	31,06 %
51-100 books per year	18,48 %	16,35 %
>100 books per year	5,32 %	6,81 %

Table 3 – Demographic Information and Reading Habits of Participants

5.2 Effects of Nudging on User Behavior

We conducted a comprehensive analysis of the potential effects of the digital nudges from multiple perspectives [2].

5.2.1 Effect on Resulting Reading Lists

We present the number of items added to the reading lists by the participants in different groups, along with the number of preferred genre and off-profile items in Table 4. Additionally, we provide the mean number of favorited items per user in Table 5, considering all items and only off-profile items.

Based on the data presented in Table 5, the results indicate that the participants in the treatment group tended to add slightly more items to their reading lists on average. However, it is worth noting that these differences were not statistically significant based on a Student's

[2] To promote reproducibility and transparency, we have made the anonymized data we collected available online at <https://github.com/gaahbie/unbook>.

	Control	Treatment	Σ
From preferred genres	1,718	1,535	**3,253**
Off-profile	751	873	**1,624**
Σ	**2,469**	**2,408**	**4,877**

Table 4 – Absolute numbers of books placed on reading list (favorited)

	Control	Treatment	*p*-value
Mean (std) of favorited items per user (total)	4.54 (1.09)	4.64 (1.01)	0.12
Mean (std) of fav'ed items per user (off-profile)	1.38 (1.10)	1.68 (1.13)	< 0.001

Table 5 – Average number of books placed on reading list (favorited)

t-test with α=0.05. It is possible that the suggestion we provided to users to favorite 3 to 5 books before using the app influenced the observed behavior.

Upon analyzing only off-profile items that participants placed on their reading list, we discovered a significant difference. The digital nudges applied to these items were effective, and participants in the treatment group placed a considerably higher number of off-profile items on their reading list. A chi-squared test was conducted with the data presented in Table 4, which revealed that the differences were statistically significant, with a p <0.001.

Furthermore, we compared the *mean* number of off-profile favorited items in the treatment and control groups. The results of a Student's t-test revealed statistical significance (p <0.001), providing further evidence that the nudges effectively encouraged users to engage with off-profile content. Completing this analysis, Table 6 shows the number of participants who placed at least one off-profile item on their reading list. The data demonstrate that the nudges stimulated more users in the treatment group to select at least one off-profile item than the control group. Again, a chi-squared test revealed statistical significance (p <0.001). These findings indicate that the nudges were effective in prompting more users to explore off-profile items in our experiment. To validate that the statistical significance was not due to an overpowered study, we repeated the analyses on three smaller subsamples (N=532, with 272 in the control group and 260 in the treatment group) selected randomly. In all cases, the results showed statistical significance with α=0.05.

	Control	Treatment	Σ
At least one off-profile item	407 (74,82 %)	440 (84,62 %)	**847**
No off-profile item	137 (25,18 %)	80 (15,38 %)	**214**
Σ	**544**	**520**	**1,064**

Table 6 – Numbers of users who placed at least one off-profile item on their reading list.

Our primary finding supports our central hypothesis that *digital nudges can be an effective means to stimulate users to explore items from genres that they were previously not among their preferred ones*. It is worth noting that participants in our study already showed a high level of openness to exploration without the nudges. In the control group, about 30 % of the reading list items were off-profile, and this number increased to approximately 36 % in the treatment group. We suggest that the significant proportion of off-profile items in the control group may, in part, be due to position or order bias, as every second item was an off-profile item.

As part of our investigation, we aimed to explore whether the nudges helped counteract popularity biases in the selected books. In many domains, we observe long-tail distributions where a small set of items receive most attention, resulting in the *rich gets richer* effect. This effect can be detrimental to both business and fairness objectives. To assess the impact of nudges on such biases, we calculated the Gini coefficient for the popularity of book genres in the treatment and control groups, inspired by the work of (JANNACH *et al.*, 2015). The Gini index ranges from zero to one and indicates the level of imbalance in data distribution, with higher values suggesting a more significant concentration.

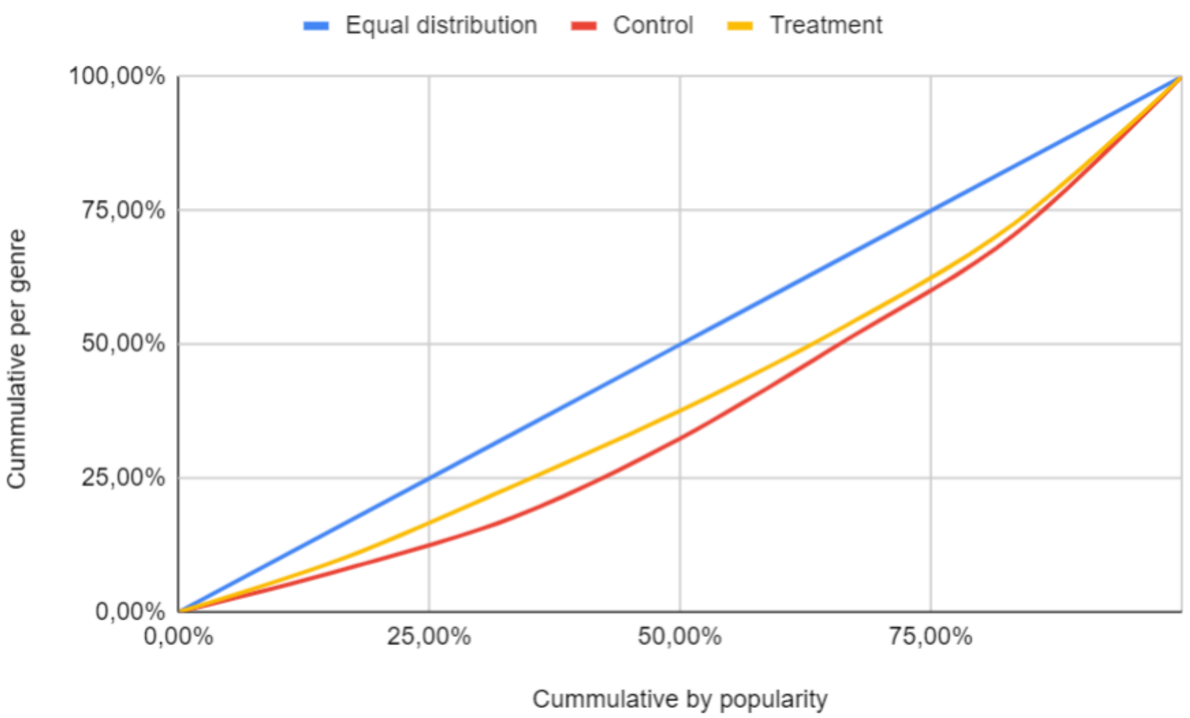

Figure 11 – Gini index of popularity of genres in reading lists

The Lorenz curve displayed in Figure 11 illustrates the distribution of genre preferences and reveals that the treatment group is closer to the line of perfect equality, suggesting a more balanced distribution with the nudges. The Gini coefficient for the treatment group is 0.17, while the control group has a higher concentration at 0.23. Our analysis indicates that nudging led to a greater selection of books from the genre "non-fiction", and a decrease in popularity for the genre "romance" in the treatment group.

5.2.2 Effect on Exploration Behavior

In addition to the changes made to the reading lists, we also collected data on other user interactions during the study. One of the most relevant statistics in this context is the number of times participants clicked the "show more" button to inspect the details of a book recommendation. The data collected is presented in Table 7.

	Control	Treatment
Inspect item details (Preferred)	439 (0.80)	503 (0.96)
Inspect item details (Off-Profile)	224 (0.41)	322 (0.62)
Avg. nb. “inspect item details” events per user	1.21	1.58

Table 7 – Summary of other recorded user interactions. Numbers in parentheses show the average number of “show more” clicks per user.

The analysis of user interactions during the study reveals some interesting findings. Specifically, participants in the treatment group showed a 30 % increase in inspecting the details of book recommendations compared to the control group (from 1.21 to 1.58 "show more" clicks). Additionally, participants in the treatment group were found to inspect more items from their preferred genres, despite there being no difference in the presentation of items compared to the control group. The statistical significance of these differences was confirmed by a chi-squared test based on the data presented in Table 7 ($p = 0.037$).

In addition to tracking participants’ item additions to their reading lists, we also monitored the frequency of item *removals*. The data revealed that there were similar numbers of item removals in both the treatment and control groups, regardless of whether the removed items were off-profile or from the participants’ preferred genres. In the control group, we observed 25 item removals, while in the treatment group, there were 31 removals. Interestingly, the differences between the two groups were primarily driven by an increase in the removal of preferred genre books in the treatment group. However, given the small absolute numbers of recorded events, we did not find any statistically significant differences between the groups. Overall, these results suggest that the nudges did not have a negative impact on participants, such as inducing them to initially add off-profile items to their reading lists only to remove them later.

We also analyzed the position in the list from *where* participants picked the books they added to their reading lists. Figure 12 presents the normalized frequency of item selections for each position in the list, for both treatment and control groups. In the control group, a clear order effect was observed for the items from the preferred genres (i.e., those with even position numbers). On the other hand, the off-profile items (i.e., those with odd position numbers) showed a clear reduction in selection frequency, as evident from the figure.

Regarding the selection of items in the treatment group, we noticed an order effect, with the neighboring items displaying smaller differences, especially at the start of the list. However, we also observed that the nudges’ impact often appeared to decrease as the list progressed. We

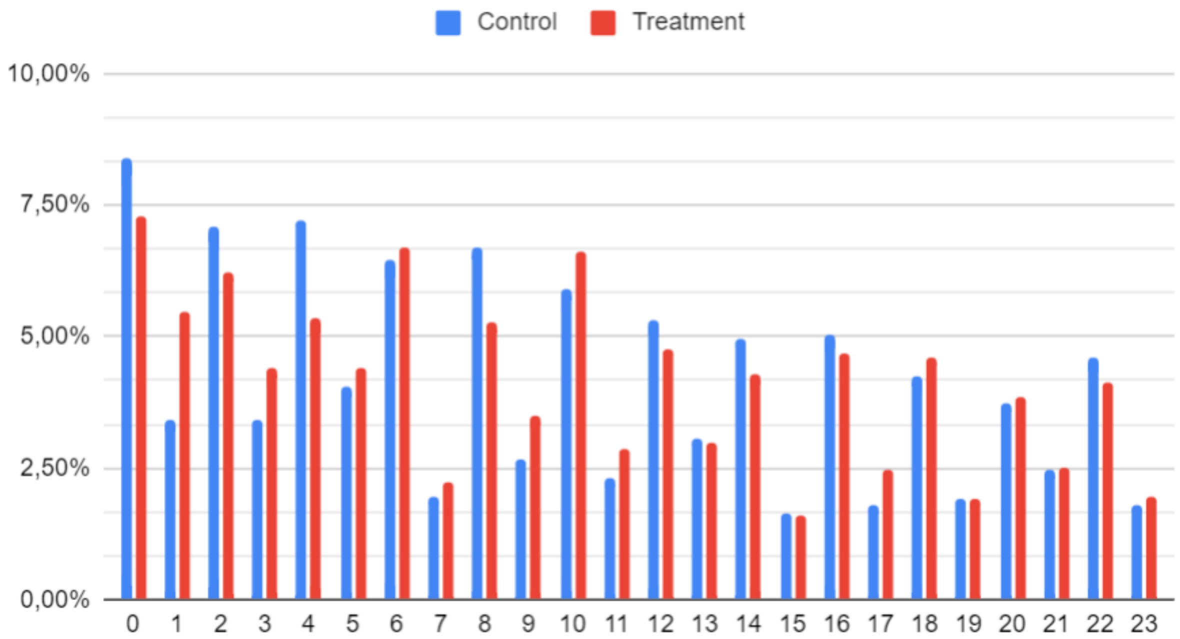

Figure 12 – Distribution of percentages of books being placed in reading lists at different list positions.

additionally monitored how far participants scrolled down the list and identified the last book they viewed, but we found no group differences. Notably, about 70 % of participants in both the treatment and control groups scrolled through the entire 24-item list.

In addition to the aforementioned analyses, we investigated the *time* taken by participants to make their selections. The average time taken by participants in the control group was 252.41 seconds, or roughly 4.2 minutes (std=187.66). On the other hand, the treatment group took an average of 331.30 seconds, or approximately 5.5 minutes, for the same task (std=218.56). This is an increase of more than 30 %, which is statistically significant (p<0.001) according to a Welch's t-test and a robust statistic analysis of the heteroscedastic data. We should note that we removed outliers that were beyond three standard deviations from the mean for this analysis. Even without removing outliers, the differences between the control and treatment groups are similar and statistically significant.

The additional time needed to select items for the reading list is expected since participants in the treatment group interacted with more items, as illustrated in Table 7. It is worth noting that various factors could cause an increase in the number of interactions and time spent with a list of recommendations. It is possible that the nudges positively affected participants' engagement, leading them to explore more options thoroughly. Alternatively, the nudges may have drawn attention to irrelevant items, resulting in unnecessary effort and distractions for the participants. To gain more insight into these questions, we will examine the findings of the post-task questionnaire in the following section.

5.2.3 Relationship between Prior Interest Diversity and Exploration

In this subsection, we conducted an analysis to determine whether the prior diversity of interests affects the effects of the nudges. One hypothesis was that participants who initially declared more preferred genres might be more easily nudged towards off-profile items because they have a predisposition to be more open. To investigate this aspect, we computed the correlation between the number of initially preferred genres and the *fraction of off-profile* items in the final reading lists. We found no such correlation ($\rho = 0.08$) across all users, and no correlation was found when considering the treatment and control groups separately. Moreover, since the average number of favorited books was not different for the treatment and control groups, no correlation existed between the number of preferred genres and the absolute number of selected off-profile items.

To investigate the potential impact of *extreme* diversity in genre preferences on the effectiveness of the nudges, we divided the participants based on the median number of declared genre preferences. The distribution of the number of declared genre preferences is shown in Figure 13. We then analyzed the data separately for two groups of participants: those who declared only one or two preferred genres and those who declared four or five preferred genres.

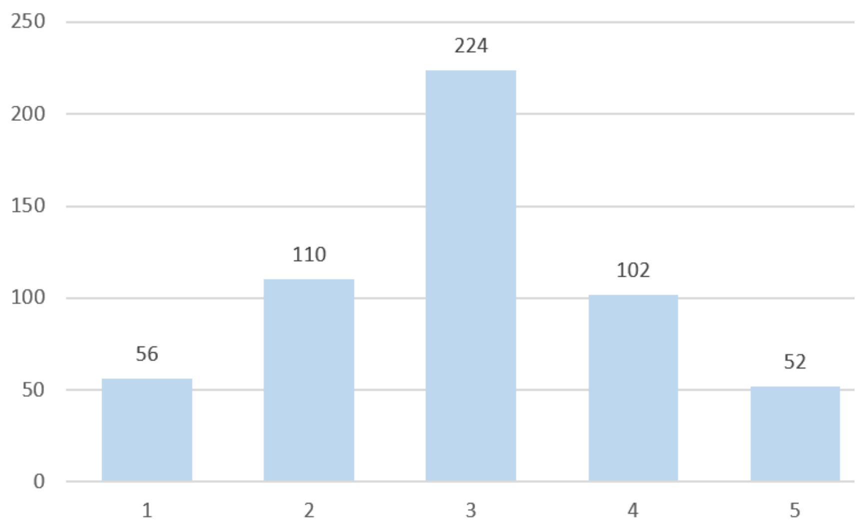

Figure 13 – Distribution of number of declared preferred genres. The x-axis shows the number of preferred genres, and the y-axis the number of participants per group.

We conducted further analyses to investigate whether participants' prior genre preferences played a role in their selection behavior regarding off-profile items. However, we did not find any significant correlation between the two variables in any of the subgroups analyzed. This was true for both participants with low and high numbers of declared preferred genres and for both the treatment and control groups. The correlation values did not exceed 0.2, which is considered very weak. Therefore, we can conclude that there is *no indication that the effectiveness of digital nudges depends on the prior preference diversity of the participants*.

5.2.4 Effects on Quality Perception and Behavioral Intentions

Table 8 presents the results of the post-task questionnaire which aimed to understand the participants' perceptions of the recommendations, their beliefs, satisfaction, and behavioral intentions. The table shows the means and standard deviations for the different groups, as well as the statistical significance determined using a Wilcoxon test.[3] In the table, p-values below 0.05 are marked with two stars, while p-values below 0.1 are marked with one star.[4]

		Control	Treatment	p-value
Q1	The items recommended to me matched my interests.	4.24 (1.04)	4.06 (1.08)	0.095*
Q2	The app provided some really good recommendations.	4.20 (0.96)	4.11 (0.94)	0.293
Q3	The items recommended to me are diverse.	4.38 (0.87)	4.31 (0.91)	0.567
Q4	The recommendations helped me discover new books which I find interesting.	3.97 (1.18)	4.05 (1.13)	0.287
Q5	I was surprised by some of the recommendations.	3.74 (1.21)	3.71 (1.16)	0.350
Q6	The amount of information available was sufficient to make the decisions.	3.87 (1.16)	3.91 (1.21)	0.525
Q7	I understood why the items were recommended to me.	4.18 (1.09)	4.03 (1.1)	0.025**
Q8	Finding books to put on my reading list was easy	4.07 (1.19)	3.89 (1.23)	0.058*
Q9	My overall satisfaction with my selection of books on my reading list is high.	4.19 (1.03)	4.05 (1.08)	0.007**
Q10	My overall satisfaction with the app is high.	4.25 (0.94)	4.10 (0.98)	0.029**
Q11	I like the app, and I would use it again in the future.	4.48 (0.87)	4.31 (0.95)	0.017**
Q12	I am eager to read the book(s) I have put on the reading list.	4.37 (0.93)	4.32 (0.97)	0.004**

Table 8 – Mean responses and standard deviations for post-task questionnaire. P-values lower than 0.05 are marked with **; p-values lower than 0.1 are marked with *.

In Table 8, we observe that the differences in the means for the 5-point response scale are generally small and below 5 %. In terms of the main recommendation quality factors, such

[3] We note that the data is not normally distributed (p<0.01) based on Shapiro-Wilk tests. It is also important to mention that we are testing 12 independent hypotheses in this analysis, based on constructs from the ResQue model (PU; CHEN; HU, 2011).

[4] We additionally performed an analysis with robust Winsorized estimates of means and standard deviations and Welch's t-tests. The significance results were well aligned with the outcomes of the Wilcoxon test.

as attractiveness, diversity, novelty, serendipity, and information sufficiency (Q2-Q6), we did not observe any statistically significant differences between the treatment and control groups. However, for Q1 (accuracy), participants in the treatment group found the recommendations to be slightly ($p<0.1$) less relevant than participants in the control group. Additionally, participants in the treatment group found it slightly ($p<0.1$) less easy to find suitable books (Q8).

Regarding the negative effects of the nudges, we found significant differences (with $p<0.05$) for the more "indirect" aspects of transparency (Q7), satisfaction (Q9, Q10), and the behavioral intentions of the participants (Q11, Q12). Specifically, we observed a general trend towards slightly worse perceptions in various dimensions in the presence of the nudges.

One noteworthy observation is that the accuracy of the recommendations (Q1) was slightly lower in the presence of the nudges. It should be noted that the recommendations were the same for all participants with the same genre preferences. One possible explanation is that the nudges drew the attention of some users who then found the off-profile items to be of little relevance to them[5]. As our results have shown, participants in the control group examined fewer off-profile items and therefore focused on items that fell into their preferred genres, resulting in an overall better perception of relevance. Surprisingly, however, the diversity perception did not increase in the treatment group and actually decreased slightly. This is unexpected as one might have assumed that the off-profile items would contribute to diversity.

In general, the small yet statistically significant decreases in various aspects may be linked to the participants' perception of relevance. This is supported by the path coefficients obtained in the validation of the ResQue's Structural Equation Model (SEM) model (PU; CHEN; HU, 2011), which showed that accuracy had a significant impact on usefulness in their study. Usefulness, in turn, influenced overall satisfaction with the recommender system, and satisfaction was positively related to the intention to use the app again in the future.

5.2.5 Relationships Between Variables

To examine if the ResQue framework's findings in (PU; CHEN; HU, 2011) are also applicable to our study, we created a path model using our post-task questionnaire data following the ResQue structure. The path coefficients resulting from the model were generated using IBM AMOS software, and the responses from all participants are illustrated in Figure 14. We found the model to be a good fit, with a comparative fit index (CFI) of 1.000 and a Root Mean Square Error of Approximation (RMSEA) of 0.000. The absence of multicollinearity between accuracy, diversity, novelty, and information sufficiency was confirmed through a collinearity analysis. To enhance readability, we have not included the covariances between these variables in Figure 14.

The results of our path model are consistent with those reported in (PU; CHEN; HU,

[5] In addition to the decrease in accuracy, the perceived transparency (Q7) also decreased significantly in the treatment group.

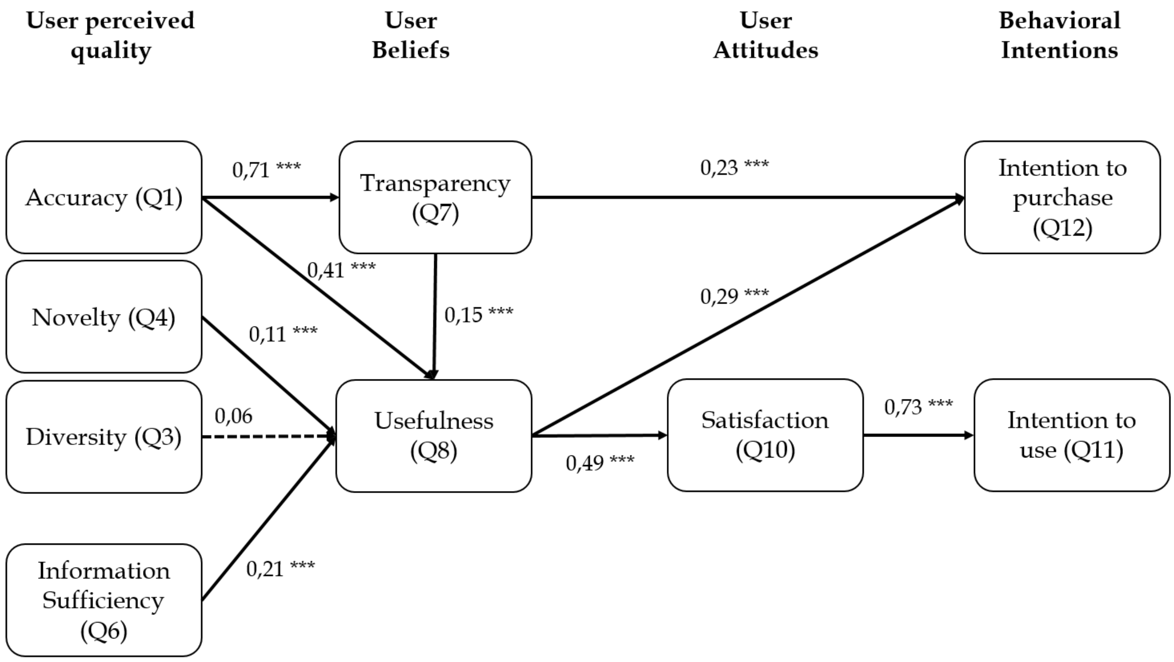

Figure 14 – Path model following the ResQue framework based on observed variables. *** indicates p-values lower than 0.001.

2011). Specifically, we found that accuracy is the most influential factor in quality perception, with a strong impact on transparency and perceived usefulness. Novelty and diversity have a smaller impact, and diversity does not significantly influence usefulness at the 0.05 level, as in (PU; CHEN; HU, 2011). Perceived usefulness, as expected, affects user satisfaction, which in turn affects users' behavioral intentions to accept recommendations and use the system in the future.

It is worth noting that not all questionnaire items were included in the path model presented in Figure 14, but only those that corresponded clearly to the constructs in the ResQue framework. Additionally, since we did not have a variable for user control in our model, we directly connected transparency with usefulness in our path model to investigate the relationship, which is indirect in the ResQue framework. In contrast to ResQue, our path model connects observed variables as we do not have latent variables based on different indicator variables. However, despite these differences, our work provides additional evidence for the validity of the ResQue model in the evaluation of recommender systems.

5.2.6 Free-form Inputs

In the final step of our analysis, we examined the free-text feedback provided by participants at the end of the experiments. Approximately 100 participants from each group provided feedback, and those comments that referred to the recommendations (as opposed to the app) were analyzed. Many users expressed a desire for more recommendations. In the control group, 16% of users mentioned that the recommended list was too short, and 6% noted that the recommendations included books from genres they did not select. The treatment group had similar results, with 15% commenting that a list of 24 books may be too short, and 5% of users stating that the recommendations had off-profile items. Overall, we did not observe any differences in the feedback given in the control and treatment groups. However, an interesting observation in the treatment group was that no comments were related to the nudges, even though they effectively influenced participants' behavior and choices.

5.3 Final Considerations

The purpose of this study was to investigate the effectiveness of digital nudges in encouraging users to explore off-profile items on their reading lists. To achieve this aim, we conducted a randomized controlled experiment comparing a treatment group, which received nudges, with a control group, which did not. Our statistical analysis revealed evidence that digital nudges can effectively encourage users to engage with off-profile content, which can help counteract popularity biases. This finding represents a significant contribution to the literature as it shows that digital nudges can stimulate users to explore genres that were previously not among their preferred ones.

Furthermore, our analysis of the data demonstrated that participants in the treatment group inspected 30% more details of book recommendations than the control group, and they explored more items from their preferred genres. However, we also found that the impact of the nudges decreased as the list progressed. We also investigated whether the prior diversity of interests affects the effectiveness of the nudges, and the results showed that there was no correlation between the number of initially preferred genres and the fraction of off-profile items in the final reading lists. Additionally, extreme diversity in genre preferences did not have a significant impact on the effectiveness of the nudges.

Finally, we examined the effects of the nudges on users' quality perceptions and behavioral intentions. Our results revealed that while the nudges had a small yet statistically significant negative impact on various dimensions of users' perceptions, such as accuracy, transparency, satisfaction, and behavioral intentions, they were still effective in promoting engagement with off-profile content.

In conclusion, this study provides new insights into the potential of digital nudges to encourage users to explore off-profile items on their reading lists. However, the study also has some limitations that need to be acknowledged. In our final chapter, we discuss these limitations and possible explanations for our findings. We also suggest directions for future research to improve our understanding of how digital nudges can be effectively used to encourage users to explore off-profile content.

CHAPTER

6

CONCLUSION

6.1 Outlook & Future Work

The study conducted in this master thesis suggests that digital nudging can be effective in encouraging users to explore content outside their typical preferences. However, our findings also suggest that excessive use of nudges may lead to negative effects on users' perceptions of quality and their future use of the system. As a result, future research should examine the impact of different levels of nudging to determine the trade-off between encouraging exploration and maintaining user satisfaction with the system.

To further advance our understanding of the impact of digital nudging on user behavior, future studies should investigate whether individual characteristics play a role in how users perceive and respond to nudges. For example, do gender or personality traits, such as openness, affect the effectiveness of nudges in encouraging users to explore content beyond their past preferences?

To explore these questions, future studies can consider two approaches. First, they can vary the level of diversity in recommendations for each user based on their individual traits, as suggested in prior research (WU; CHEN; HE, 2013; GUO *et al.*, 2020). This would involve adjusting the number of off-profile items presented to the user. Second, researchers can tailor the type of nudge used based on the user's personality traits, as done in previous studies (GUO *et al.*, 2020). These investigations can provide valuable insights into the individual factors that influence the effectiveness of digital nudging and how they can be used to create personalized recommendation systems.

The present study examined the impact of applied nudges on the bookmarking behavior of participants in the domain of book recommendations. However, in other domains, such as health and nutrition, prior research suggests that not all types of nudges are equally effective (BERGER; MÜLLER; NÜSKE, 2020; JESSE; JANNACH; GULA, 2021). While it would be interesting to analyze the effectiveness of different nudges in our application setting, such an in-depth analysis was beyond the scope of our current research. Moreover, our study design cannot draw definitive conclusions about the effectiveness of individual nudges due to confounding factors such as position effects. Preliminary analysis suggests that the non-hybrid Nudge N4 (i.e., "Social reference point") may be less effective than other nudges. However, a more comprehensive analysis is needed to confirm these indications. Future studies may address this gap in knowledge and provide a more nuanced understanding of the effectiveness of different nudges in book recommendation systems.

Our study provides valuable insights into the potential of digital nudges to increase literacy in a country. This is a relatively unexplored area, and our findings could contribute to achieving this important societal goal. However, there are many other areas where nudging could be useful in promoting positive societal outcomes. For instance, nudges could be used to encourage online news readers to consume articles with opposing viewpoints, which may help to reduce filter bubbles and prevent radicalization. Future studies are needed to explore the effectiveness of nudging in these and other areas.

In the context of news recommendations, a crucial question is how different off-profile recommendations should be from a user's past preferences. In our experiment, we considered all non-preferred genres as equally distant from the preferred ones. However, for political or controversial topics, it may be advisable to avoid off-profile recommendations that are too far removed from the reader's past tendencies. The theory of the Overton Window of Political Possibility (LEHMAN, 2010) suggests establishing a spectrum of "acceptable" opinions and then selecting off-profile content for nudging that has a reasonable chance of being considered by readers. Moreover, as discussed in Vermeulen (2022), users may want more control over their level of exploration. This area could be a fruitful area for future research.

6.2 Threats to Validity

This section contributes to the understanding of the study's ecological validity and generalizability and their implications for the findings' validity and applicability.

6.2.1 *Ecological Validity*

Our study aimed to maintain ecological validity by involving real users of a book recommendation platform. However, like other user studies, we acknowledge potential limitations in terms of realism. One such limitation could arise if participants are not genuinely interested and engaged when interacting with the system, especially if their participation is voluntary and unpaid. To address this, we observed high ratings from participants on the usefulness of the app and their intention to use it in the future. Furthermore, we found that participants spent an average of five minutes adding items to their reading lists, and over 50% inspected at least one item detail page, indicating genuine engagement. Additionally, almost 80% of participants scrolled to the end of the 24-item list, suggesting that they took the task seriously. Therefore, we believe that the ecological validity of our study is high. However, it is essential to note that our study focused on changes in behavior in terms of adding items to a reading list, and we cannot infer whether participants actually bought, read, or enjoyed the selected books.

6.2.2 Generalizability

Our study focused on book recommendations, and therefore we cannot confidently conclude whether our findings would apply to other domains.

It's important to note that our participants were real users of a book recommendation platform, making them an essential subset of the user population for book recommender systems. However, our study participants may not represent the average population of large online bookstores, such as Amazon.com, which offers a vast assortment of books. On average, our participants read only a few dozen books per year, and they may be more open to exploring new titles than more occasional readers or those with limited interests. Additionally, our study participants were mostly female, and further research is necessary to determine if gender differences impact the effectiveness of nudges.

Nonetheless, our findings are consistent with prior research demonstrating the effectiveness of nudges in other areas such as healthy food choices, energy-saving, and movie recommendations (as discussed in Chapter 3). As a result, our study contributes to the body of knowledge on digital nudging. However, the effectiveness of specific nudges in different settings remains a critical area for future exploration.

BIBLIOGRAPHY

AKIN, L. K. **Information overload. A multi-disciplinary explication and citation ranking within three selected disciplines: Library studies, psychology/psychiatry, and consumer science, 1960-1996**. TX, USA: Texas Woman's University, 1997. Citation on page 23.

BAEZA-YATES, R. Bias on the web. **Communications of the ACM**, v. 61, p. 54–61, 05 2018. Citation on page 24.

BAEZA-YATES, R.; RIBEIRO-NETO, B. A. Modern information retrieval - the concepts and technology behind search, second edition. In: . New York, USA: Addison-Wesley, 2011. Citation on page 51.

BERGER, M.; MÜLLER, C.; NÜSKE, N. Digital nudging in online grocery stores: Towards ecologically sustainable nutrition. 2020. Citations on pages 26, 54, and 71.

BURKE, R.; ROBIN. Hybrid web recommender systems. **LNCS**, v. 4321, 01 2007. Citation on page 30.

CARABAN, A.; KARAPANOS, E.; GONcALVES, D.; CAMPOS, P. 23 ways to nudge: A review of technology-mediated nudging in human-computer interaction. p. 1—15, 2019. Citations on pages 26 and 45.

CASTELLS, P.; VARGAS, S.; WANG, J. Novelty and diversity metrics for recommender systems: Choice, discovery and relevance. 2011. Citation on page 44.

CELMA Òscar. Music recommendation and discovery - the long tail, long fail, and long play in the digital music space. Springer, 2020. Citation on page 25.

CHEN, J.; DONG, H.; WANG, X. lei; FENG, F.; WANG, M.-C.; HE, X. Bias and debias in recommender system: A survey and future directions. **ArXiv**, abs/2010.03240, 2020. Citations on pages 24, 31, and 33.

CHEN., L.; PU., P. Trust building in recommender agents. **Proceedings of the 1st International Workshop on Web Personalisation, Recommender Systems and Intelligent User Interfaces - WPRSIUI, (ICETE 2005)**, SciTePress, p. 135–145, 2005. Citation on page 47.

CHEN, L.; YANG, Y.; WANG, N.; YANG, K.; YUAN, Q. How Serendipity Improves User Satisfaction with Recommendations? A Large-Scale User Evaluation. **The World Wide Web Conference**, 2019. Citation on page 44.

CHOPRA, K.; WALLACE, W. A. Trust in electronic environments. **36th Annual Hawaii International Conference on System Sciences, 2003. Proceedings of the**, p. 10 pp.–, 2003. Citation on page 35.

Christopher Westland, J. Lower bounds on sample size in structural equation modeling. **Electronic Commerce Research and Applications**, v. 9, n. 6, p. 476–487, 2010. Citation on page 60.

DELDJOO, Y.; JANNACH, D.; BELLOGÍN, A.; DIFONZO, A.; ZANZONELLI, D. Fairness in recommender systems: Research landscape and future directions. 2022. Citation on page 33.

DIAS, M. B.; LOCHER, D.; LI, M.; EL-DEREDY, W.; LISBOA, P. J. The value of personalised recommender systems to e-business: A case study. p. 291–294, 2008. Citation on page 43.

Efthimios Bothos; Dimitris Apostolou; Gregoris Mentzas. A recommender for persuasive messages in route planning applications. **2016 7th International Conference on Information, Intelligence, Systems and Applications (IISA)**, p. 1–5, 2016. Citation on page 47.

EKSTRAND, M. D.; HARPER, F. M.; WILLEMSEN, M. C.; KONSTAN, J. A. User perception of differences in recommender algorithms. p. 161–168, 2014. Citation on page 25.

ELSWEILER, D.; TRATTNER, C.; HARVEY, M. Exploiting food choice biases for healthier recipe recommendation. p. 575–584, 2017. Citations on pages 31 and 46.

EPPLER, M. J.; MENGIS, J. The concept of information overload: a review of literature from organization science, accounting, marketing, mis, and related disciplines. **IEEE Engineering Management Review**, v. 38, p. 3–3, 2004. Citation on page 23.

FOGG, B. Persuasive technology: Using computers to change what we think and do. **Ubiquity**, v. 2002, n. December, p. 2, 2002. Citation on page 46.

GE, M.; GEDIKLI, F.; JANNACH, D. Placing high-diversity items in top-n recommendation lists. 2011. Citation on page 25.

GE, M.; JANNACH, D.; GEDIKLI, F.; HEPP, M. Effects of the placement of diverse items in recommendation lists. **ICEIS 2012 - Proceedings of the 14th International Conference on Enterprise Information Systems**, v. 2, 2012. Citations on pages 23 and 25.

GELFERT, A. Fake news: A definition. **Informal Logic**, v. 38, p. 84–117, 2018. Citation on page 33.

GENA, C.; GRILLO, P.; LIETO, A.; MATTUTINO, C.; VERNERO, F. When personalization is not an option: An in-the-wild study on persuasive news recommendation. **Information**, v. 10, n. 10, p. 300, 2019. Citations on pages 47 and 52.

GUO, Y.; ZHANG, Z.; GUO, Y.; GUO, X. Nudging personalized password policies by understanding users' personality. **Computers & Security**, v. 94, p. 101801, 2020. Citation on page 71.

HANSEN, C.; MEHROTRA, R.; HANSEN, C.; BROST, B.; MAYSTRE, L.; LALMAS, M. Shifting consumption towards diverse content on music streaming platforms. **Proceedings of the 14th ACM International Conference on Web Search and Data Mining**, p. 238–246, 2021. Citation on page 45.

HEITZ, L.; LISCHKA, J. A.; BIRRER, A.; PAUDEL, B.; TOLMEIJER, S.; LAUGWITZ, L.; BERNSTEIN, A. Benefits of diverse news recommendations for democracy: A user study. **Digital Journalism**, forthcoming, p. 1–21, 2022. Citation on page 45.

JACOBY, J.; SPELLER, D. E.; BERNING, C. K. Brand choice behavior as a function of information load: Replication and extension. **Journal of Consumer Research**, v. 1, p. 33–42, 1974. Citation on page 23.

JANNACH, D.; LERCHE, L.; KAMEHKHOSH, I.; JUGOVAC, M. What recommenders recommend: an analysis of recommendation biases and possible countermeasures. **User Modeling and User-Adapted Interaction**, Springer Nature, v. 25, n. 5, p. 427–491, 2015. Citation on page 62.

JANNACH, D.; ZANKER, M.; FELFERNIG, A.; FRIEDRICH, G. Recommender systems: An introduction. **Recommender Systems: An Introduction**, 01 2010. Citations on pages 23, 24, and 29.

JESSE, M.; JANNACH, D. Digital nudging with recommender systems: Survey and future directions. **Computers and Human Behavior Reports**, v. 3, 2021. Citations on pages 17, 37, 38, 39, 43, 46, 47, 52, 53, and 54.

JESSE, M.; JANNACH, D.; GULA, B. Digital nudging for online food choices. **Frontiers in Psychology**, v. 12, 2021. Citations on pages 26, 46, 53, 54, and 71.

JOACHIMS, T. Evaluating retrieval performance using clickthrough data. p. 79–96, 2002. Citation on page 51.

KAMEHKHOSH, I.; BONNIN, G.; JANNACH, D. Effects of recommendations on the playlist creation behavior of users. **User Modeling and User-Adapted Interaction**, v. 30, p. 285–322, 2019. Citation on page 43.

KAPOOR, K.; KUMAR, V.; TERVEEN, L.; KONSTAN, J. A.; SCHRATER, P. “I like to Explore Sometimes”: Adapting to Dynamic User Novelty Preferences. **Proceedings of the 9th ACM Conference on Recommender Systems**, p. 19–26, 2015. Citation on page 44.

KIM, J.-K.; CHOI, I. Y.; LI, Q. Customer satisfaction of recommender system: Examining accuracy and diversity in several types of recommendation approaches. **Sustainability**, v. 13, p. 6165, 2021. Citation on page 48.

KONSTAN, J. A.; MCNEE, S. M.; ZIEGLER, C.-N.; TORRES, R.; KAPOOR, N.; RIEDL, J. Lessons on applying automated recommender systems to information-seeking tasks. In: **AAAI**. Sao Paulo – SP: Novatec Editora, 2006. Citation on page 35.

KONSTAN, J. A.; RIEDL, J. Recommender systems: from algorithms to user experience. **User Modeling and User-Adapted Interaction**, v. 22, p. 101–123, 2011. Citation on page 34.

LAWRENCE, R. D.; ALMASI, G. S.; KOTLYAR, V.; VIVEROS, M. S.; DURI, S. S. Personalization of supermarket product recommendations. **Data Min. Knowl. Discov.**, v. 5, n. 1–2, p. 11–32, 2001. Citation on page 43.

LEHMAN, J. G. **An Introduction to the Overton Window of Political Possibility**. 2010. Mackinac Center for Public Policy, <https://www.mackinac.org/12481>. Citation on page 72.

LEX, E.; KOWALD, D.; SEITLINGER, P.; TRAN, T. N. T.; FELFERNIG, A.; SCHEDL, M. Psychology-informed recommender systems. **Found. Trends Inf. Retr.**, v. 15, p. 134–242, 2021. Citation on page 37.

LI, J.; LU, K.; HUANG, Z.; SHEN, H. T. Two birds one stone: on both cold-start and long-tail recommendation. p. 898–906, 2017. Citation on page 25.

LIANG, Y.; WILLEMSEN, M. C. The role of preference consistency, defaults and musical expertise in users' exploration behavior in a genre exploration recommender. **Fifteenth ACM Conference on Recommender Systems**, 2021. Citation on page 46.

LOPS, P.; GEMMIS, M. de; SEMERARO, G. Content-based recommender systems: State of the art and trends. In: RICCI, F.; ROKACH, L.; SHAPIRA, B.; KANTOR, P. (Ed.). **Recommender Systems Handbook**. New York: Springer, 2011. p. 73–105. Citation on page 43.

LU, J.; WU, D.; MAO, M.; WANG, W.; ZHANG, G. Recommender system application developments: A survey. **Decis. Support Syst.**, v. 74, p. 12–32, 2015. Citation on page 24.

LUDEWIG, M.; LATIFI, S.; MAURO, N.; JANNACH, D. Empirical analysis of session-based recommendation algorithms. **User Modeling and User-Adapted Interaction**, v. 31, n. 1, p. 149–181, 2021. Citations on pages 25 and 44.

LUNARDI, G. M.; MACHADO, G. M.; MARAN, V.; OLIVEIRA, J. P. M. de. A metric for filter bubble measurement in recommender algorithms considering the news domain. **Appl. Soft Comput.**, v. 97, p. 106771, 2020. Citation on page 34.

MALHOTRA, N. K. Information load and consumer decision making. **Journal of Consumer Research**, v. 8, p. 419–430, 1982. Citation on page 23.

MARTIN, R. C. **Clean Architecture**. <https://blog.cleancoder.com/uncle-bob/2012/08/13/the-clean-architecture.html>. Accessed: 2022-04-28. Citation on page 56.

MCINERNEY, J.; LACKER, B.; HANSEN, S.; HIGLEY, K.; BOUCHARD, H.; GRUSON, A.; MEHROTRA, R. Explore, exploit, and explain: Personalizing explainable recommendations with bandits. **Proceedings of the 12th ACM Conference on Recommender Systems**, p. 31–39, 2018. Citation on page 44.

MCNEE, S. M.; RIEDL, J.; KONSTAN, J. A. Being accurate is not enough: How accuracy metrics have hurt recommender systems. **CHI '06 Extended Abstracts on Human Factors in Computing Systems**, p. 1097–1101, 2006. Citation on page 44.

MESKE, C.; POTTHOFF, T. The dinu-model – a process model for the design of nudges. In: **ECIS**. Guimarães, Portugal: In Proceedings of the 25th European Conference on Information Systems (ECIS), 2017. Citation on page 38.

MILANO, S.; TADDEO, M.; FLORIDI, L. Recommender systems and their ethical challenges. **AI & SOCIETY**, p. 1–11, 2020. Citation on page 24.

MIRSCH, T.; LEHRER, C.; JUNG, R. Digital nudging: Altering user behavior in digital environments. **Wirtschaftsinformatik und Angewandte Informatik**, p. 634–648, 2017. Citation on page 52.

MOLS, F.; HASLAM, S. A.; JETTEN, J.; STEFFENS, N. K. Why a nudge is not enough: A social identity critique of governance by stealth. **European Journal of Political Research**, v. 54, n. 1, p. 81–98, 2015. Citation on page 46.

MUNSON, S.; RESNICK, P. Presenting diverse political opinions: how and how much. **Proceedings of the SIGCHI Conference on Human Factors in Computing Systems**, 2010. Citation on page 34.

NANOU, T.; LEKAKOS, G.; FOUSKAS, K. G. The effects of recommendations' presentation on persuasion and satisfaction in a movie recommender system. **Multimedia Systems**, v. 16, p. 219–230, 2010. Citations on pages 35, 36, and 37.

NGUYEN, T. T.; HARPER, F. M.; TERVEEN, L. G.; KONSTAN, J. A. User personality and user satisfaction with recommender systems. **Information Systems Frontiers**, v. 20, p. 1173–1189, 2018. Citation on page 48.

NGUYEN, T. T.; HUI, P.-M.; HARPER, F. M.; TERVEEN, L. G.; KONSTAN, J. A. Exploring the filter bubble: the effect of using recommender systems on content diversity. **Proceedings of the 23rd international conference on World wide web**, 2014. Citation on page 33.

OOSTERHUIS, H. Learning from user interactions with rankings: A unification of the field. **ArXiv**, abs/2012.06576, 2020. Citation on page 51.

PARISER, E. **The filter bubble : what the Internet is hiding from you**. New York: Penguin Press, 2011. Citations on pages 33 and 34.

PASSE, J.; DRAKE, C.; MAYGER, L. K. Homophily, echo chambers, & selective exposure in social networks: What should civic educators do? **The Journal of Social Studies Research**, 2018. Citation on page 34.

PU, P.; CHEN, L.; HU, R. A user-centric evaluation framework for recommender systems. **Proceedings of the 5th ACM Conference on Recommender Systems**, p. 157–164, 2011. Citations on pages 15, 54, 55, 66, 67, and 68.

RENAUD, K.; ZIMMERMANN, V. Nudging folks towards stronger password choices: providing certainty is the key. **Behavioural Public Policy**, Cambridge University Press, v. 3, n. 2, p. 228–258, 2019. Citation on page 54.

RESNICK, P.; VARIAN, H. R. Recommender systems. **Commun. ACM**, v. 40, p. 56–58, 1997. Citation on page 23.

RICCI, F.; ROKACH, L.; SHAPIRA, B. Introduction to recommender systems handbook. In: **Recommender Systems Handbook**. New York: Springer, 2011. p. 1–35. ISBN 978-0-387-85820-3. Citation on page 29.

RICCI, F.; ROKACH, L.; SHAPIRA, B.; KANTOR, P. B. Recommender systems handbook. In: **Recommender Systems Handbook**. New York: Springer, 2010. p. 1–35. ISBN 978-0-387-85820-3. Citation on page 35.

RIEGER, A.; DRAWS, T.; THEUNE, M.; TINTAREV, N. This item might reinforce your opinion: Obfuscation and labeling of search results to mitigate confirmation bias. **Proceedings of the 32nd ACM Conference on Hypertext and Social Media**, 2021. Citation on page 47.

SCHELENZ, L. Diversity-aware Recommendations for Social Justice? Exploring User Diversity and Fairness in Recommender Systems. **UMAP 2021 - Adjunct Publication of the 29th ACM Conference on User Modeling, Adaptation and Personalization**, p. 404–410, 2021. Citation on page 25.

SCHNEIDER, C.; WEINMANN, M.; BROCKE, J. vom. Digital nudging: Guiding online user choices through interface design. **Communications of the ACM**, v. 61, n. 7, p. 67–73, 2018. Citations on pages 26 and 45.

STARKE, A. D.; WILLEMSEN, M. C.; SNIJDERS, C. With a little help from my peers: Depicting social norms in a recommender interface to promote energy conservation. **Proceedings of the 25th International Conference on Intelligent User Interfaces**, p. 568–578, 2020. Citation on page 46.

STARKE, A. D.; WILLEMSEN, M. C.; TRATTNER, C. Nudging healthy choices in food search through visual attractiveness. **Frontiers in Artificial Intelligence**, v. 4, 2021. Citation on page 46.

THALER, R.; SUNSTEIN, C. **NUDGE: Improving Decisions About Health, Wealth, and Happiness**. New York: Penguin, 2009. ISBN 9780141040011. Citations on pages 25, 37, 45, and 52.

TVERSKY, A.; KAHNEMAN, D. Judgment under uncertainty: Heuristics and biases. **Science**, v. 185, p. 1124 – 1131, 1974. Citation on page 37.

VERMA, A.; WADHWA, A.; SINGH, N.; BENIWAL, S.; KAUSHAL, R.; KUMARAGURU, P. Followee management: Helping users follow the right users on online social media. **IEEE/ACM International Conference on Advances in Social Networks Analysis and Mining (ASONAM)**, p. 1286–1290, 2018. Citation on page 47.

VERMEULEN, J. To nudge or not to nudge: News recommendation as a tool to achieve online media pluralism. **Digital Journalism**, 2022. Citations on pages 47 and 72.

VICARIO, M. D.; BESSI, A.; ZOLLO, F.; PETRONI, F.; SCALA, A.; CALDARELLI, G.; STANLEY, H. E.; QUATTROCIOCCHI, W. The spreading of misinformation online. **Proceedings of the National Academy of Sciences**, v. 113, p. 554 – 559, 2016. Citation on page 34.

WANG, Y.; MA, W.; ZHANG, M.; LIU, Y.; MA, S. A survey on the fairness of recommender systems. **ACM Journal of the ACM (JACM)**, 2022. Citation on page 32.

WEINMANN, M.; SCHNEIDER, C.; BROCKE, J. vom. Digital nudging. **Business & Information Systems Engineering**, v. 58, p. 433–436, 2016. Citation on page 37.

WERNECK, H.; SANTOS, R.; SILVA, N.; PEREIRA, A. C.; MOURÃO, F.; ROCHA, L. Effective and diverse poi recommendations through complementary diversification models. **Expert Systems with Applications**, Elsevier, v. 175, p. 114775, 2021. Citation on page 47.

WU, W.; CHEN, L.; HE, L. Using personality to adjust diversity in recommender systems. **Proceedings of the 24th ACM Conference on Hypertext and Social Media**, p. 225–229, 2013. Citation on page 71.

YOO, K. H.; GRETZEL, U. Creating more credible and persuasive recommender systems: The influence of source characteristics on recommender system evaluations. In: RICCI, F.; ROKACH, L.; SHAPIRA, B.; KANTOR, P. B. (Ed.). **Recommender Systems Handbook**. New York: Springer, 2011. p. 455–477. ISBN 978-0-387-85820-3. Citation on page 37.

YOO, K. H.; GRETZEL, U.; ZANKER, M. **Persuasive Recommender Systems - Conceptual Background and Implications**. New York: Springer, 2013. i-vi, 1-59 p. (Springer Briefs in Electrical and Computer Engineering). ISBN 978-1-4614-4702-3. Citations on pages 37, 46, and 47.

ZIEGLER, C.-N.; MCNEE, S. M.; KONSTAN, J. A.; LAUSEN, G. Improving recommendation lists through topic diversification. **Proceedings of the 14th International Conference on World Wide Web**, p. 22–32, 2005. Citation on page 44.

APPENDIX

A

FIGURES & TABLES

Figures 15 – 18 show stylized and translated screenshots of the app (treatment group) with the nudges N1 to N4 applied. Tables 9 – 14 then list the books that were recommended for each genre and which nudges were applied.

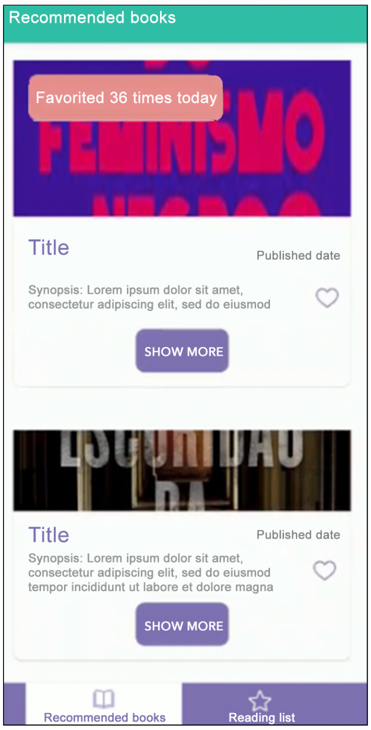

Figure 15 – Screen Capture (Treatment Group), with nudge N1 applied.

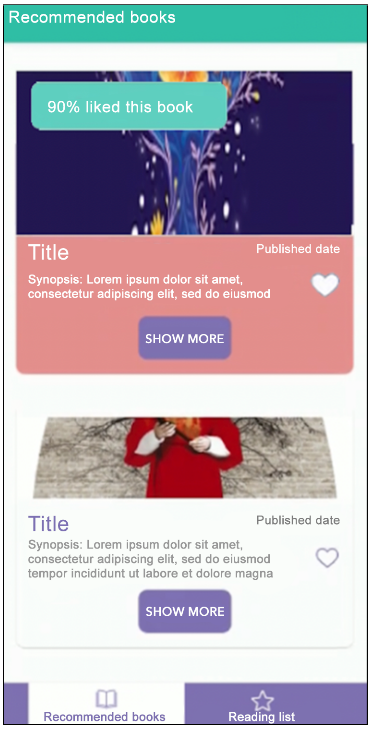

Figure 16 – Screen Capture (Treatment Group), with nudge N2 applied.

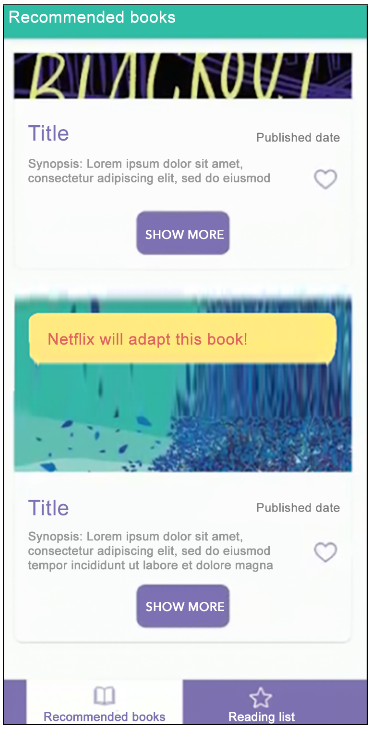

Figure 17 – Screen Capture (Treatment Group), with nudge N3 applied.

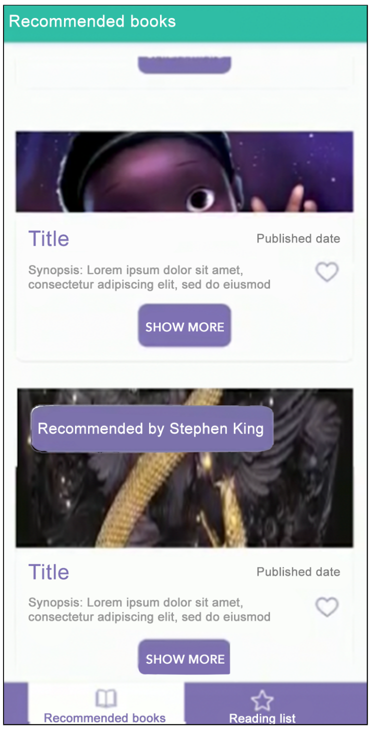

Figure 18 – Screen Capture (Treatment Group), with nudge N4 applied.

Title	Author	N1	N2	N3	N4
Como o Rei de Elfhame aprendeu a odiar histórias	Holly Black				X
Estilhaça-me: 1	Tahereh Mafi				X
O guia definitivo do mochileiro das galáxias	Douglas Adams				X
A Princesa Prometida	William Goldman			X	
O Conto da Aia	Margaret Atwood			X	
Os Pilares da Terra	Ken Follett			X	
1984: Edição com Postais + Marcador	George Orwell		X		
Sombra e Ossos	Leigh Bardugo		X		
Trono de vidro (Vol. 1)	Sarah J. Maas		X		
A maldição do mar	Shea Ernshaw	X			
O canto mais escuro da floresta	Holly Black	X			
O Labirinto Do Fauno	Guillermo Del Toro	X			
A Fonte	CS Luis				
Box A Arma Escarlate: 4	Renata Ventura				
Depois	Stephen King				
Em Algum Lugar nas Estrelas	Clare Vanderpool				
Fábulas árabes	M. M. Jarouche				
Harry Potter e a Pedra Filosofal	J.K. Rowling				
Kindred: laços de sangue	Octavia E. Butler				
Maldição	Marie O'Regan				
MARVEL - OS PRIMEIROS 80 ANOS	Culturama				
Necronomicon em Netvilly	Kleber Inácio Da Silva				
Os reis do Wyld	Nicholas Eames				
Trilogia da Fundação - Deluxe	Isaac Asimov				
Vingadores: Guerra Infinita -	Steve Behling				

Table 9 – List of *fantasy* books and applied nudges.

Title	Author	N1	N2	N3	N4
Drácula - Dark Edition	Bram Stoker				X
Nunca Saia Sozinho	Charlie Donlea				X
O Retrato de Dorian Gray	Oscar Wilde				X
Joyland	Stephen King			X	
Na Escuridão da Mente	Paul Tremblay			X	
Tempo Estranho	Joe Hill			X	
A estrada da noite – Edição Luxo	Joe Hill		X		
Box Terríveis Mestres	Edgar Allan Poe		X		
It: A coisa	Stephen King		X		
Bom dia, você está morto!	Chell Sant'Ana	X			
Carrie (Com brindes)	Stephen King	X			
Frankenstein: O clássico está vivo!	Mary Shelley	X			
A canção de Bêlit: a tigresa e o leão					
A ilha do tesouro	Robert Louis Stevenson				
A pequena caixa de Gwendy	Stephen King				
Box HP Lovecraft	Howard Phillips Lovecraft				
Calmaria Forçada	Rosane Montalvão				
Edgar Allan Poe - Medo Clássico	Edgar Allan Poe				
Eu sei o que vocês fizeram no verão passado	Lois Duncan				
Frankenstein, ou o Prometeu Moderno	M W Shelley				
O bosque das coisas perdidas	Shea Ernshaw				
O Colecionador	John Fowles				
O médico e o monstro	Robert Louis Stevenson				
Round 6 - por dentro da série	Park Minjoon				
Tempo Estranho	Joe Hill				
Tumular: a sete palmos do inferno	Thunder Dellú				
Uma mulher na escuridão	Charlie Donlea				

Table 10 – List of *horror* books and applied nudges.

Title	Author	N1	N2	N3	N4
Ame pessoas, use coisas	Joshua Fields Millburn				X
Ed & Lorraine Warren: Vidas Eternas	Robert Curran				X
Monstros	Simon Sebag Montefiore				X
A Bailarina de Auschwitz	Edith Eva Eger			X	
Persópolis	Marjane Satrapi			X	
Sapiens - Uma Breve História da Humanidade	Yuval Noah Harari			X	
Box Memórias da Segunda Guerra Mundial	Winston Churchill		X		
Serial Killers - Anatomia do Mal	Harold Schrechter		X		
Ted Bundy: Um Estranho ao Meu Lado	Ann Rule		X		
ARQUIVOS SERIAL KILLERS	Ilana Casoy	X			
Enquanto eu respirar	Ana Michelle Soares	X			
Lady Killers	Tori Telfer	X			
50 Cent - Minha História, Minha Verdade	50 Cent				
A marca da vitória	Phil Knight				
Alpha Girls	Julian Guthrie				
BTK Profile: Máscara da Maldade	Roy Wenzl				
Decolonialidade e pensamento afrodiaspórico	Joaze Bernardino-Costa				
Denali	Ben Moon				
Elvis Presley.	Gillian G. Gaar				
Escritos de Uma Vida	Sueli Carneiro				
Jovens heróis da União Soviética	Alex Halberstadt				
Meus primeiros 21	Nikki Sixx				
Na estrada com os Ramones	Monte A. Melnick				
O homem mais rico da Babilônia	George Samuel Clason				
Por que escrever?	Philip Roth				
Quem tem medo do feminismo negro?	Djamila Ribeiro				
Rebelde - Autobiografia do Criador de Conan	Robert E. Howard				
Stalin: Uma biografia	Robert Service				
Teologia do Corpo	São João Paulo II				
The Beatles.	Vários Autores				
Tolkien e a Grande Guerra	John Garth				

Table 11 – List of *non-fiction* books and applied nudges.

Title	Author	N1	N2	N3	N4
A vida invisível de Addie LaRue	V.E. Schwab				X
Amor & Gelato	Jenna Evans Welch				X
O morro dos ventos uivantes	Emily Brontë				X
Americanah	Chimamanda Ngozi Adichie			X	
Daisy Jones & The Six	Taylor Jenkins Reid			X	
O Rouxinol	Kristin Hannah			X	
É Assim que Acaba	Colleen Hoover		X		
Os Bridgertons, um amor de famíli	Julia Quinn		X		
Sem julgamentos	Meg Cabot		X		
História é tudo que me deixou	Adam Silvera	X			
Pássaro e serpente (Vol. 1)	Shelby Mahurin	X			
Teto Para Dois	Beth O'leary	X			
novembro, 9	Colleen Hoover				
A soma de todos os afetos	Fabíola Simões				
Até o verão terminar	Colleen Hoover				
Em outra vida, talvez?	Taylor Jenkins Reid				
Malibu renasce	Taylor Jenkins Reid				
Mil beijos de garoto	Tillie Cole				
Missão romance	Lyssa Kay Adams				
Mr. 365	Ruth Clampett				
O acordo	Elle Kennedy				
O conde enfeitiçado – Edição Luxo	Julia Quinn				
O lado feio do amor	Colleen Hoover				
O Palácio de Papel	Miranda Cowley Heller				
Quase Uma Família	Sherryl Woods				
Rosas Esquecidas	Martha Hall Kelly				
Todas as suas (im)perfeições	Colleen Hoover				
Todo esse tempo	Rachael Lippincott				
Um pai para o meu bebê?	Tamires Barcellos				

Table 12 – List of *romance* books and applied nudges.

Title	Author	N1	N2	N3	N4
Deixada Para Trás	Charlie Donlea				X
Gente Ansiosa	Fredrik Backman				X
O último julgamento	Scott Turow				X
Amada	Toni Morrison			X	
Mestre das Chamas	Joe Hill			X	
O Chamado do Cuco	Robert Galbraith			X	
Mestre das Chamas	Joe Hill		X		
Um de nós está mentindo	Karen M. McManus		X		
Verity	Colleen Hoover		X		
Assassinato no Expresso do Oriente	Agatha Christie	X			
Billy Summers	Stephen King	X			
O Homem de Giz	C. J. Tudor	X			
A Contrapartida - Livro 2: O Contra-ataque	Uranio Bonoldi				
A garota na neve	Danya Kukafka				
A Lista de Convidados	Lucy Foley				
A paciente silenciosa	Alex Michaelides				
A Salvação: Uma história de Vampiro	Diego Sousa				
A Última Festa	Lucy Foley				
Box Sherlock Holmes	Arthur Conan Doyle				
Coleção Agatha Christie - Box 1	Agatha Christie				
Detalhe final (Myron Bolitar – Livro 6)	Harlan Coben				
E não sobrou nenhum	Agatha Christie				
Em fogo lento	Paula Hawkins				
Estado de alerta	David Klass				
Mentiras incendiárias	Jennifer Lynn Alvarez				
Não Confie em Ninguém	Charlie Donlea				
O Clube do Crime das Quintas-Feiras	Richard Osman				
O mistério de Agatha Christie	Benedict Benedict				
Um de nós é o próximo	Karen M. McManus				
Um Segredo em Provence	Walter Barbosa				
Veu de Veronica, O	Raphael Prats				

Table 13 – List of *suspense* books and applied nudges.

Title	Author	N1	N2	N3	N4
Conectadas	Clara Alves				X
O príncipe cruel (Vol. 1 O Povo do Ar)	Holly Black				X
Por lugares incríveis	Jennifer Niven				X
O Livro da Selva	Rudyard Kipling			X	
O Ódio que Você Semeia	Angie Thomas			X	
Sulwe	Lupita Nyong'o			X	
A rainha vermelha	Victoria Aveyard		X		
A Seleção: 1	Kiera Cass		X		
Vermelho, branco e sangue azul	Casey McQuiston		X		
Coraline	Neil Gaiman	X			
Heartstopper: Dois garotos, um encontro (vol. 1)	Alice Oseman	X			
Última parada	Casey McQuiston	X			
A rainha do nada (Vol. 3 O Povo do Ar)	Holly Black				
Aristóteles e Dante...	Benjamin Alire Sáenz				
Ash	Malinda Lo				
Assim você me mata	Karen McManus				
Blackout: O amor também brilha no escuro	Dhonielle Clayton				
Bruxa Natural	Arin Murphy-Hiscock				
Confusão é meu nome do meio	Stephanie Tromly				
Coragem	Raina Telgemeier				
Corte de chamas prateadas	Sarah J. Maas				
Decifra-me	Tahereh Mafi				
Em fogo alto - com brinde exclusivo	Elizabeth Acevedo				
Fat Chance: A vez de Charlie Vega	Crystal Maldonado				
Harry Potter e a pedra filosofal	J.K. Rowling				
Metamorfose	Franz Kafka				
Os últimos jovens da Terra	Max Brallier				
Um dia a alma transborda	Marina Leme				
Unifica-me + brindes	Tahereh Mafi				

Table 14 – List of books for *young adults* and applied nudges.

ICMC USP
SÃO CARLOS